The International Journal of Biostatistics

Volume 7, Issue 1

2011

Article 39

Dose-Finding Designs: The Role of Convergence Properties

Assaf P. Oron, *University of Washington*

David Azriel, *The Hebrew University of Jerusalem*

Peter D. Hoff, *University of Washington*

Recommended Citation:

Oron, Assaf P.; Azriel, David; and Hoff, Peter D. (2011) "Dose-Finding Designs: The Role of Convergence Properties," *The International Journal of Biostatistics*: Vol. 7: Iss. 1, Article 39.

DOI: 10.2202/1557-4679.1298

Dose-Finding Designs: The Role of Convergence Properties

Assaf P. Oron, David Azriel, and Peter D. Hoff

Abstract

It is common for novel dose-finding designs to be presented without a study of their convergence properties. In this article we suggest that examination of convergence is a necessary quality check for dose-finding designs. We present a new convergence proof for a nonparametric family of methods called “interval designs,” under certain conditions on the toxicity-frequency function F . We compare these conditions with the convergence conditions for the popular CRM one-parameter Phase I cancer design, via an innovative numerical sensitivity study generating a diverse sample of dose-toxicity scenarios. Only a small fraction of scenarios meet the Shen-O’Quigley convergence conditions for CRM. Conditions for “interval design” convergence are met more often, but still less than half the time. In the discussion, we illustrate how convergence properties and limitations help provide insight about small-sample behavior.

KEYWORDS: adaptive designs, continual reassessment method, dose finding, phase I clinical trials, cumulative cohort design, up-and-down, item-response theory, sensory studies

Author Notes: All authors thank the editors and anonymous referees for their comments and suggestions, which have helped improve this manuscript substantially. Oron would like to thank Nancy Flournoy for her ongoing mentorship on dose-finding designs, and for introducing him to Azriel. Azriel would like to thank his dissertation advisors, Micha Mandel and Yosef Rinott.

1 Introduction

1.1 Motivation

A wide range of experiments in medicine, science and engineering can be modeled statistically as percentile-finding (more commonly known as “dose-finding”). Treatments vary by the value of a single continuous variable t , and responses are dichotomized to, say, “yes” and “no”. The experiment’s goal is to find the treatment that would generate a “yes” response for a given fraction p of the population. Under common statistical assumptions this is equivalent to finding $F^{-1}(p)$ – the $100p$ -th percentile of $F(t)$, the cumulative distribution function of response thresholds in the sampled population. In this article we limit ourselves to the discrete-treatment case, in which treatments are drawn from a fixed set of predetermined levels $\Omega \equiv \{d_1, d_2, \dots, d_K : d_1 < d_2 < \dots < d_K\}$. Following decades of scant methodological activity on dose finding, it has attracted considerable attention recently via the Phase I cancer clinical trial application (hereafter simply “Phase I”). Phase I researchers typically seek the maximum tolerated dose (MTD): the dose level in Ω that is closest to the specified target toxic-response frequency p . This task might be more accurately described as “dose selection”; however, we will use the more commonly encountered term “dose finding.”

The most widely discussed approach in the new wave of Phase I designs uses repeated estimation of F to guide dose allocation decisions (O’Quigley et al., 1990). This approach is called by O’Quigley and Zohar (2006) “designs with memory”. We will use the term “long-memory designs” – as opposed to “short-memory designs”, such as the family of methods known as up-and-down (U&D, Dixon and Mood, 1948). Short-memory designs do *not* use an estimation procedure for treatment allocation.

Most recent Phase I methodology discussions have revolved around heuristic arguments and simulation studies rather than around proving a design’s asymptotic properties. Apparently, the rationale is that Phase I is a small-sample application – 10 to 50 subjects – and therefore convergence matters little. This mindset is misguided. The common thread among dose finding designs of all types is the use of a self-correction mechanism to concentrate treatments around target, thus improving the precision of its detection. If the self-correction mechanism is sound, then treatment allocations should eventually converge to some stationary behavior with desirable properties. Conversely, if convergence cannot be guaranteed under realistic conditions, then the self-correction mechanism itself might be of limited utility. In other words, convergence should be viewed as a *necessary quality criterion* for dose finding designs (albeit not a sufficient one, since small-sample behavior does need to be examined closely as well).

1.2 Currently Available Convergence Results

We should emphasize that the term “convergence” in the dose finding context refers to **allocation convergence**, i.e., the convergence of *the distribution of allocated treatments* to some asymptotic distribution with known properties (rather than, say, convergence of F estimates to the true F). Among dose finding designs, the convergence behavior of the U&D family of short-memory designs is the best understood. U&D designs generate Markov chains, converging to a stationary random walk whose dose-allocation distribution is centered close to target (Derman, 1957, Durham and Flournoy, 1995, Gezmu, 1996, Gezmu and Flournoy, 2006, Oron and Hoff, 2009). Some researchers see the random-walk spread of allocations around target as a drawback of U&D, because a secondary goal of clinical trials is to provide trial volunteers with a treatment as optimal as feasible under study limitations. In contrast to U&D, most long-memory designs currently available, and certainly the most well-known ones such as the continual reassessment method (CRM, O’Quigley et al., 1990), promise to eventually concentrate treatments solely at the MTD.

There are very few published results on the convergence of long-memory dose finding designs. The best-known among them is Shen and O’Quigley (1996)’s proof for designs with one-parameter models. This proof is often quoted as evidence that the most popular long-memory design – one-parameter CRM – converges under general model misspecification. However, Cheung and Chappell (2002) demonstrated that the Shen-O’Quigley proof requires rather strong restrictions on F . They conjectured that the restrictions might be partially relaxed. The conjecture has yet to be proven; both sets of conditions will be revisited below in Section 4. Two other proofs related to two-parameter long-memory designs (Zacks et al., 1998, Roy et al., 2009) – escalation with overdose control (EWOC Babb et al., 1998) and D -optimal Bayesian Phase I designs (Haines et al., 2003), respectively – require that the model correctly specify F . It is generally accepted that Phase I researchers seldom possess information sufficient to specify the form of F .

Azriel et al. (2011) have recently placed a general limitation on long-memory convergence: they proved that designs aiming to allocate treatments solely at the MTD *cannot* be strongly consistent for an arbitrary dose-response curve F . The proof applies to both parametric and nonparametric designs. Given this knowledge, our attention in convergence studies should turn towards learning more about the convergence conditions of various long-memory designs, and comparing them in terms of simplicity and restrictions on F .

This article presents a convergence proof, under straightforward conditions on F , for a recently introduced class of nonparametric long-memory designs called here “interval designs” (Yuan and Chappell, 2004, Ivanova et al., 2007). Prelimi-

nary terms and results to aid the proof appear in the next section, followed by the convergence theorem itself in Section 3. For the reader's convenience, the theorem's proof and the proof of an important lemma leading to it, as well as some additional definitions, are relegated to an appendix. Section 3 will also briefly describe a counterexample disproving convergence of another nonparametric long-memory approach we call "point designs." We revisit the Shen and O'Quigley (1996) proof and Cheung and Chappell (2002) conjecture for one-parameter designs in Section 4, comparing the restrictiveness of the conditions for one-parameter designs and interval designs. The article ends with a discussion examining the results, contrasting them with short-memory convergence, and also suggesting how convergence properties might be leveraged to better understand small-sample behavior.

2 Preliminaries

2.1 The Designs

We describe the dose finding problem via a latent-variable model, and use the Phase I terms "dose" and "toxicity" interchangeably with "treatment", "response", respectively. Let $Y(t) \sim \text{Bernoulli}(F(t))$ be a binary toxicity response of some dose strength t , with the toxicity-rate function $F(t)$ strictly monotone increasing but not directly observable. As mentioned in the introduction, rather than precisely estimate the target, researchers are often content with identifying the MTD, i.e., the dose level closest to target according to some distance criterion; we denote the MTD as $d_{j^*} \in \Omega$, and define it as $j^* \equiv \operatorname{argmin}_{1 \leq k \leq K} |F(d_k) - p|$, the most common criterion used in practice.

Consider a sequential design treating $m_i \geq 1$ subjects at cohort i , $i = 1, 2, \dots$. For simplicity and without loss of generality with respect to our proofs, from here on we assume that all cohorts are of size 1, and index successive treatments as the r.v.'s T_i , $i = 1, \dots, n, \dots$, with $t_i \in \Omega \equiv \{d_1, d_2, \dots, d_K : d_1 < d_2 < \dots < d_K\}$. The toxicity responses Y_i are assumed independent given the T_i . We define sequential designs rather broadly as any design allocating cohort i based on the outcomes of cohorts $i = 1, \dots, i-1$, and some pre-specified set of rules that usually involve design parameters, both fixed and data-estimable.

All *long-memory* designs use the raw toxicity frequencies, which can be seen as Binomial point estimates of F given n observations,

$$\hat{F}_n(d_k) \equiv \frac{\sum_{i=1}^n y_i \mathbf{1}[T_i = d_k]}{\sum_{i=1}^n \mathbf{1}[T_i = d_k]} \quad \forall d_k \in (t_1, \dots, t_n), \quad (1)$$

Figure 1: Illustration of the designs discussed in the article, using data from the Mathew et al. (2004) Phase I experiment targeting 30% toxicity. Shown is the situation after Cohort 3, administered at 35 mg/m²/week. Observed toxicity rates are marked with 'X'. CRM (left) would allocate to Cohort 4 the dose closest to where the posterior model curve (solid line) crosses the dashed horizontal $y = 0.3$ line, i.e., 25 mg/m²/week. The point design (center) would follow the same principle, but using a nonparametric curve obtained via isotonic-regression interpolation; hence, Cohort 4 would receive 30 mg/m²/week. The interval design (right) looks only at the toxicity rate at the dose administered to Cohort 3. Since this rate falls above the interval marked by two dashed horizontal lines, the allocation will de-escalate one level to 30 mg/m²/week.

where y_i is the binary toxicity outcome (0 or 1). Parametric designs use the $\hat{F}$ values indirectly as inputs to fit a family of model curves $G(t, \theta)$ approximating the true curve F , via estimation of θ . Nonparametric long-memory designs use the $\hat{F}$ directly, with a possible modification to ensure monotonicity via standard methods (e.g., Robertson et al., 1988). Following is the definition of the “interval design”, and also of another nonparametric design called here “point design.”

Definition 1 (i) An “interval-based nonparametric long-memory” Phase I design (hereafter, “interval design”) starts at an arbitrary dose. At each subsequent step, supposing d_k is the currently-administered dose, then if

$$\hat{F}(d_k) \in (p - \Delta p_1, p + \Delta p_2),$$

with $\Delta p_1 > 0$, $\Delta p_2 > 0$ predetermined constants, then d_k will be administered again. If $\hat{F}(d_k) \leq p - \Delta p_1$, d_{k+1} will be administered (unless $k = K$ in which case d_K will be administered again), and vice versa if $\hat{F}(d_k) \geq p + \Delta p_2$.

(ii) A “point-based nonparametric long-memory” Phase I design (hereafter, “point design”) starts at an arbitrary dose. At each subsequent step, the design allocates the next cohort to the level whose (possibly monotonicized) $\hat{F}(d_k)$ value is closest to p . If there are several levels with the same $\hat{F}(d_k)$, the lowest of them will be allocated.

The point design was suggested by Leung and Wang (2001); it is a direct variation on parametric designs such as CRM (Figure 1, left), with the parametric curve $G(\hat{\theta})$ replaced by a monotone nonparametric interpolation of F between dose levels (Figure 1, center). The interval design's principle is different; one might call it “narrow long-memory” since the allocation decision is based on prior outcomes at the current dose only (Figure 1, right). Rather than look for some optimal dose at each cohort, the allocation would repeat the existing dose as long as it is deemed “close enough” to the true MTD, i.e., within interval boundaries. Different versions of the interval design were put forth by Yuan and Chappell (2004) and Ivanova et al. (2007), the latter introducing the acronym “cumulative cohort design” (CCD). The interval design does not allow for skipping dose levels between consecutive cohorts.¹

2.2 Convergence of Raw Point Estimates

The following result might seem self-evident, but surprisingly we could not find any published reference where it appears. All long-memory proofs mentioned above implicitly use it or some analogues to it as an intermediate step, but the result is general and useful enough to be presented alone.

Lemma 1 *Using the terminology introduced above, for any dose finding design and all $d_k \in \Omega$, $\hat{F}_n(d_k) \rightarrow F(d_k)$ almost surely as $n_k \rightarrow \infty$, where n_k is the number of patients to whom d_k was assigned.*

In other words: the convergence of raw point estimates to their expectations, which is guaranteed for i.i.d. sampling by the Laws of Large Numbers, also holds for sequential dose finding designs. The proof, using martingale theory, is in the appendix.

¹Interestingly, Ivanova et al. (2007) present their design as related to the U&D family, mainly because their interval width recommendations were based upon a numerical exhaustive search over U&D designs. However, since the dose allocation themselves are based upon estimates of F rather than upon recent cohort outcomes as in U&D, the design is in fact long-memory.

3 Theoretical Convergence Results

3.1 Interval-Design Convergence

Theorem 1 (i) *Dose allocations in interval designs converge almost surely to d_{j^*} , if the latter maintains*

$$F(d_{j^*}) \in (p - \Delta p_1, p + \Delta p_2), \quad (2)$$

and if d_{j^} is also the only level satisfying*

$$F(d_{j^*}) \in [p - \Delta p_1, p + \Delta p_2]. \quad (3)$$

(ii) Almost-sure convergence to d_{j^} will also occur if $F(d_1) \geq p + \Delta p_2$ (meaning that $j^* = 1$) or $F(d_K) \leq p - \Delta p_1$ (meaning that $j^* = K$).*

The theorem's proof is in the Appendix. It relies upon the point-convergence result of Lemma 1: once enough information is accumulated at any given treatment level, it becomes clear whether this level's true F value lies within the interval or outside it. In the former case, the design's rules mandate that allocations will remain at the same level, which according to the theorem's conditions is the MTD. In the latter case, allocations will eventually move away from the current level towards the MTD. The proof's logic immediately leads to the following results when conditions are violated:

Corollary 1 (i) *If no dose level satisfies (2) but $p \in [F(d_1), F(d_K)]$, an interval design would eventually oscillate almost surely between the two doses whose F values straddle the target interval.*

(ii) If there are multiple levels satisfying (2), an interval design will converge almost surely to one of these levels. However, convergence to d_{j^} itself (i.e., to the level closest to target) is not guaranteed.*

3.2 Point-Design Convergence

The point design has a positive probability of not converging, even for rather ordinary response curves. We show this via a simple, yet generic counterexample: Assume that $p < 1/2$, $F(d_{j^*}) = p$ and $0 < F(d_{j^*-1}) < p$ (all other levels matter little). The experiment starts, as dose finding trials often do, from the lowest level. Sooner or later d_{j^*-1} is reached, and with high probability within a few cohorts we will have $\hat{F}(d_{j^*-1}) < p$, mandating escalation to d_{j^*} . Now, suppose that the

very first cohort at d_{j^*} is all toxicities; clearly the probability for that occurring is positive. Then $\hat{F}(d_{j^*}) = 1$. Since $p < 1/2$, regardless of the value of $\hat{F}(d_{j^*-1})$, it is now closer to target than $\hat{F}(d_{j^*})$ and it will be assigned. Moreover, if no dose higher than d_{j^*} is sampled (clearly, this can occur with positive probability under a suitable starting rule) then since $\hat{F}(d_{j^*}) = 1$ monotonicity at d_{j^*} will not be violated, meaning that no monotonizing corrections can modify this point estimate, which will remain too far from p for the remainder of the experiment. Hence d_{j^*} will never be assigned again. A similar argument was made by Cheung (2002). A design recently suggested by Azriel et al. (2011) can be seen as a variation on the point design, with some added randomization to prevent the phenomenon illustrated in the example above. It was proven to converge to the MTD in probability, but not almost surely.

4 Numerical Sensitivity Study

4.1 Convergence of One-Parameter Designs

How restrictive are the conditions outlined in Theorem 1 for interval designs? We now compare them with the Shen and O’Quigley (1996) conditions for one-parameter design convergence.² When properly calibrated, for any true monotone F curve a one-parameter family of CDFs can contain K curves, each of which matches F on at least one $d_k \in \Omega$, $1 \leq k \leq K$, with a monotone relationship between k and the respective values of θ . A set of auxiliary conditions in Shen and O’Quigley (1996) guarantees this calibration.

Additionally, restrictions on F were required for the proof. For a given F consider the specific parameter value θ_k such that $G(d_k | \theta_k) = F(d_k)$. Then the level which, according to $G(t | \theta_k)$ appears to be the MTD, will be called the level “nominated” by d_k under F , since it will be allocated whenever $\hat{G}$ precisely matches the true F at d_k . The crucial and most restrictive Shen-O’Quigley condition is that *under F , all levels in Ω must nominate d_{j^*} , the true MTD.*

Cheung and Chappell (2002), in their interpretation of the proof, opine that this requires a very close match between G and F along the entire dose range. They suggest that perhaps the Shen-O’Quigley (S-O) conditions were too restrictive: it might be enough for the MTD to nominate itself, for doses below the MTD to nominate higher doses than themselves, and vice versa. Thus, dose allocation might eventually be “funneled” towards the MTD. Conversely, there are scenario types

²Hereafter we will occasionally refer to that convergence result as “CRM convergence”, even though it is in fact a proof for the convergence of an analogous frequentist design.

under which a one-parameter design *cannot* be guaranteed to converge: when non-MTD levels nominate themselves, or when the MTD fails to nominate itself.

4.2 Comparing Interval-Design and One-Parameter Convergence

We numerically explored the relative restrictedness of one-parameter and interval-design convergence conditions. Since both can be directly determined from the values of F on Ω together with design-specific parameters, there is no need to simulate actual experiments. Rather, we simulated various scenarios of F on $K = 5$ and $K = 10$ dose sets, and examined whether the CRM and the interval design convergence conditions are met for each scenario. We chose the target rate $p = 0.3$, the value most commonly used in Phase I. For this target, developers of the CCD interval design recommend the interval $(0.2, 0.4)$ when $K = 6$. Result with the narrower interval $(0.25, 0.35)$ are available in Supplementary Materials. Hereafter we refer to the interval design in this simulation as “CCD”. For CRM, we used the recently popular “power” model, in which $G(d_1, \dots, d_K; \theta) = (p_1, \dots, p_K)^{\exp(\theta)}$, with the p ’s being prior toxicity rates assigned to each dose. A choice commonly encountered in the field resembles a geometrically-increasing sequence, e.g. $(p_1, \dots, p_K) = (0.05, 0.1, 0.2, 0.4, 0.8)$ for $K = 5$ (Pisters et al., 2004). Results with a sigmoid-shaped curve are available in Supplementary Materials.

In order to generate reasonably realistic scenarios without limiting ourselves to a given distribution family, and also in order to minimize the direct impact of arbitrary conscious choice upon F , we simulated increments of F in each scenario as a random Dirichlet vector. Dirichlet distribution parameters control the likelihood of generating various curves; these parameters themselves were randomly drawn out of a finite pool, producing a range of diverse, yet reasonably realistic F curves, which should be relevant for the practical dose finding problem as defined here. Additionally, lower and upper bounds were placed on increments of F to exclude scenarios in which adjacent-dose toxicity rates are virtually indistinguishable, or spaced too far apart. Figure 2 shows a random sample of 20 scenarios (out of 2500 used in the study) for each of $K = 5$ and $K = 10$. Scenarios were simulated and convergence evaluated using R (R Development Core Team, 2011). Additional details appear in Supplementary Materials.

For CCD, we distinguish between the three possible convergence outcomes proven in Theorem 1 and Corollary 1:

1. Convergence to the MTD guaranteed (rows marked “Yes”). The MTD is the only level in the interval, or the target is below/above the design dose range.
2. More than one level in the interval (rows marked “No: 2+”). Hence, only convergence to within the interval is assured, but not to the MTD itself.

Figure 2: Simulated dose-toxicity curves. Random samples of 20 out of the 2500 simulated scenarios, for $K = 5$ (left) and $K = 10$ (right).

3. No level in the interval (rows marked “No: 0”). An asymptotic oscillating behavior is expected.

For CRM, we distinguish between four possible outcomes:

1. Shen-O’Quigley conditions met (rows marked “Yes”): all levels nominate the MTD.
2. Only Cheung-Chappell conditions are met (rows marked “Cheung-Chappell”).
3. Convergence not guaranteed (rows marked “No: 2+”): multiple self-nominating levels.
4. Convergence essentially impossible (rows marked “No: 0”): MTD does not self-nominate.

We noticed a clear distinction in performance between scenarios with true MTD on the boundaries, and ones with interior MTDs, with the former being more amenable to convergence. Therefore, Table 1 reports interior-MTD (top) and boundary-MTD (bottom) scenarios separately. With interior MTDs, the full S-O conditions for CRM convergence are met only very rarely, especially when $K = 10$. On the other hand, the weaker C-C conditions are met in a quarter to half of the scenarios, and are met more often with more design levels. Observing the CCD results, exact convergence to the MTD is guaranteed in a third of the cases with $K = 5$, but far less often with $K = 10$. However, together with the multiple-level (“No: 2+”) cases, in the vast majority of simulated scenarios CCD is guaranteed to converge to within the pre-specified interval. The standard interval width of $p \pm 0.1$ seems too wide for $K = 10$.

Table 1: Comparative convergence summary of CRM and CCD designs, for a diverse ensemble of numerically-generated scenarios. All numbers in the table are in percents. Row and column labels are explained in the text.

Interior MTD Scenarios	5 Levels		10 Levels	
	CCD	CRM	CCD	CRM
	<i>N</i> = 2251		<i>N</i> = 2377	
Yes	35.5	0.7	6.2	0.0
Cheung-Chappell	–	25.9	–	43.9
No: 2+	57.6	57.2	93.8	43.0
No: 0	6.9	16.2	0.0	13.1

Boundary MTD Scenarios	CCD	CRM	CCD	CRM
	<i>N</i> = 249		<i>N</i> = 123	
Yes	42.6	10.8	22.0	16.3
Cheung-Chappell	–	16.5	–	54.5
No: 2+	43.0	48.6	62.6	10.6
No: 0	14.5	24.1	15.4	18.7

5 Discussion

5.1 Simulation Concept

We tried to minimize the effect of direct human choice on the simulated scenarios. Preliminary attempts at this type of simulation were made by Paoletti et al. (2004) and Ivanova et al. (2007), the former starting near target and choosing F increments around it out of a Normal distribution, and the latter choosing F values out of an ordered uniform distribution. We believe our method further expands the horizons for an extensive multi-scenario simulation study, and does succeed in sampling a sizable region of the space of distributions that would be considered realistic by researchers in the field. It might serve as a template for future benchmark comparative performance simulations between designs, of the type common in fields such as machine learning.

5.2 Convergence: Interval Designs and One-Parameter Designs

The article’s main mathematical result, Theorem 1, specifies the conditions for convergence of the nonparametric long-memory “interval design.” This design has yet

to enter extensive use. Our theorem should not be seen as an advocacy for or against interval designs, but rather – in line with the article’s main *conceptual* message that convergence properties do matter – as a resource to inform consultants and researchers who might consider using interval designs. As a referee pointed out, the interval design in its CCD formulation has been presented as closely related to the U&D family, and therefore some might expect it to exhibit random walk behavior. However, in essence (as shown in Section 2) CCD is a long-memory design, exhibiting the long-memory tendency to zoom in exclusively on a single dose.

If one assumes that one-parameter design convergence can be proven under the more relaxed Cheung-Chappell conditions rather than the restrictive Shen-O’Quigley conditions, then according to our numerical study neither CCD nor CRM hold an overwhelming edge in terms of prevalence of convergence.³ However, the interval design conditions are simpler and more tractable. When exclusively-converging scenarios are combined with the case of multiple levels inside the interval (marked “2+” in Table 1), then in the vast majority of simulated scenarios a CCD experiment would be guaranteed to converge to a level within a pre-specified tolerance interval of response frequencies.

As mentioned in Section 4, each design’s tuning parameters used in our numerical study were the ones recommended by method developers, rather than chosen by us in order to optimize convergence prospects. Additional simulations carried out using design parameters that we see as more convergence-friendly for each design can be found in Supplementary Materials. For example, one-parameter convergence results with $K = 5$ improved somewhat when using a sigmoid prior rather than a convex one: 32.3% of interior-MTD scenarios meet the S-O or C-C conditions compared to 26.6% with the convex prior. A more substantial improvement is seen for CCD when using a narrower interval of $(0.25, 0.35)$ with $K = 10$, suggesting that interval width should be a function of dose spacing. Based on Theorem 1, researchers using the interval design should aim to capture approximately one dose level in the interval; erring towards more than one level is probably more desirable than capturing none. In the absence of prior scientific knowledge about the slope of F around target, a total interval width of $1/K$ (i.e., $p \pm 1/2K$) should do as a rough guideline. Another practical recommendation arising from Theorem 1 is to keep the interval symmetric. Otherwise, one might encounter cases when the MTD is outside the interval while a non-MTD level is inside the interval; clearly, in such cases theorem conditions are violated and allocations will converge to the non-MTD level. The Yuan and Chappell (2004) design uses asymmetric intervals.

Our numerical study helps illuminate the recent result of Azriel et al. (2011): namely, that no design attempting to generate a sequence of exclusive MTD-only

³As this article goes to press, a proof has been found; a related manuscript is in review.

allocations, can achieve this goal almost surely for a general F . One might argue that this result notwithstanding, the actual prevalence of non-converging scenarios is very low. Table 1 indicates the opposite: for each of two very different exclusive-MTD long memory designs (CRM and CCD), convergence conditions are met in only a minority of scenarios. A second line of argument is that even when MTD convergence is not *guaranteed* for a given scenario, it is still the most common outcome in actual runs. For example, Shen and O’Quigley (1996) reported that in their numerical trials, even for non-converging scenarios the majority of simulated experimental runs did seem to settle on the MTD when allowed to continue long enough. These scenarios might have met the C-C conditions, or even have multiple self-nominating levels. It is plausible that many actual experiments would still converge correctly in the latter case. However, under scenarios for which the one-parameter design had a non-self-nominating MTD, which comprised 10% – 30% of our simulation ensemble, there is zero probability of CRM converging exclusively to the MTD, even in a single run.

5.3 Lemma 1 and Small-Sample Behavior

Lemma 1, proving the convergence of raw point estimates of F for general sequential designs, exemplifies the connection between convergence and small-sample behavior. When unlimited information accumulates around target, the true F values in the region of interest are precisely known. We can then directly determine whether the conditions for convergence are met, and witness the asymptotic behavior. In their present configuration, long-memory designs (parametric or nonparametric) operate from the start upon the implicit assumption, or rather hope, that point estimates are *already* precise and accurate, and therefore such knowledge is at hand. Unfortunately, during the small-sample stage, which includes the entire actual Phase I experiment, the Binomial point estimates are still very imprecise.

Lemma 1 provides another important piece of information: the convergence rate for long-memory designs is, at best, root- n . Thus, these designs employ a self-correction that is not only approximate (in the sense that exclusive MTD-allocation convergence cannot be guaranteed), but also rather slow considering the sample size. To further complicate matters, long-memory experiments and simulation runs often seem to settle rather early on the same dose for several consecutive cohorts, producing the mistaken notion that they have already converged. This misleading phenomenon is in fact the side effect of the slow, root- n self-correction rate.

Long-memory design convergence properties provide an interesting contrast with those of the short-memory up-and-down designs mentioned in the Introduction. The latter converge at a geometric rate: their short memory facilitates a very

quick self-correction mechanism. However, the self-correction is blunt, and asymptotic behavior meanders around target, typically spreading the bulk of allocations over 2 – 4 levels.⁴ The tradeoff at the current state of design development is between two approximate self-correction mechanisms: one of them fast, guaranteed to converge, but blunt; and the other one slower, not-quite-guaranteed, but potentially sharper and employing more versatile mechanisms. A combined design that leverages the advantages of both while avoiding the worst of either, might be a good compromise (Oron, 2007, Ch. 5).

A final note: as both referees have kindly reminded us, and as we state in the introduction, Phase I is far from being the only application where percentile-finding methods are used, and perhaps the perspective currently dominating Phase I methodology is a trifle too narrow. At least two fields – sensory studies and educational testing design – have preceded Phase I in the development of adaptive model-based methods (Watson and Pelli, 1979, Lord, 1980), and are encountering similar challenges with their theory and application. For example, Chang and Ying (2009) recently explored convergence issues for educational testing designs.

Unfortunately, since the early 1960's the flow of information and ideas between fields using dose-finding designs has been extremely limited. Statisticians are best positioned to close this gap and to encourage better synergy across fields. For example, while the Phase I field is predominantly interested in dose *selection* for Phase II, most other fields prefer an *estimate* of the target percentile. This latter goal is generally more informative, and its performance assessment can be more nuanced. Keeping the different perspectives in mind might help us develop designs that are more well-rounded for use in a variety of fields, and eventually also more useful specifically for Phase I studies.

⁴In spite of the blunt allocation distribution, up-and-down *estimates* do become sharper with time, since they rely on all the gathered information.

Appendix: Proofs

Proof of Lemma 1

In the framework introduced in the article's body, the sample space of sequential-design allocations is the space of all permissible infinite sequences of assigned doses. It is a subset of Ω^∞ , subject to the constraint of no dose skipping. Each design induces a probability distribution on sequences in this sample space; the probability of individual allocations can be calculated with knowledge of the design's rules and of F . Almost sure convergence to the MTD means that sequences ending with infinite and uninterrupted repetitions of d_{j^*} have a combined probability of 1.

On this sample space, define the random set

$$\mathbb{S} \equiv \{k : n_k \rightarrow \infty \text{ as } n \rightarrow \infty\}, \quad (4)$$

where n_k is the number of subjects assigned to d_k . In words, $\mathbb{S}$ is the set of indices for levels appearing an infinite number of times in the sequence. Obviously $\mathbb{S}$ is nonempty for all sequences in the sample space. Moreover, since the interval design does not allow for dose skipping $\mathbb{S}$ must be connected, i.e., composed of consecutive levels. Thus, the value of $\mathbb{S}$ for different sequences in the sample space can be described via an ordered pair of integer random variables $S_1 \leq S_2$: $\mathbb{S} = S_1, \dots, S_2$.

With respect to our proof, the possible configurations of $\mathbb{S}$ can be partitioned into three major subspaces A, B, C :

- A : $S_1 = S_2 = j^*$,
- B : $S_1 < S_2$ and $j^* \in \mathbb{S}$,
- C : $j^* \notin \mathbb{S}$.

Almost sure convergence to the MTD is equivalent to $\Pr(A) = 1$.

By definition of $\mathbb{S}$ we know that for all $k \in \mathbb{S}$, $n_k \rightarrow \infty$ as $n \rightarrow \infty$. The point estimates can be written as

$$\hat{F}_n(d_k) = F(d_k) + \frac{1}{n_k} \sum_{i=1}^n I(T_i = d_k)(Y_i - F(d_k)). \quad (5)$$

Now, $M_n \equiv \sum_{i=1}^n I(T_i = d_k)(Y_i - F(d_k))$ is a square integrable martingale with respect to the filtration $\mathcal{F}_n \equiv \sigma(T_1, Y_1, \dots, T_n, Y_n)$. Its quadratic variation is:

$$\sum_{i=1}^n [I(T_i = d_k)]^2 \cdot F(d_k) \cdot [1 - F(d_k)] \propto \sum_{i=1}^n I(T_i = d_k) = n_k.$$

Therefore, due to the strong law of martingales (Shirayev, 1996, p. 519, theorem 4)

$$\frac{1}{n_k} \sum_{i=1}^n I(T_i = d_k)(Y_i - F(d_k)) \rightarrow 0 \text{ a.s. } \forall k \in \mathbb{S}.$$

The lemma's statement, i.e., $\hat{F}_n(d_k) \rightarrow F(d_k) \forall k \in \mathbb{S}$, immediately follows. $\square$

Proof of Theorem 1

(i) Using the same terminology, we first establish that $\Pr(C) = 0$, which is equivalent to $\Pr(j^* \in \mathbb{S}) = 1$. We do it by contradiction, assuming w.l.o.g. that there is some specific level $s_1 > j^*$ for which $\Pr(S_1 = s_1) > 0$. In words, we assume that there are sequences with a combined positive probability of occurring, in which beyond a certain point *only levels above the MTD* are visited. From the theorem's conditions we know that $F(d_{s_1}) > p + \Delta p_2$. Due to Lemma 1, this means that for n large enough and all sequences described by the assumed event,

$$\Pr\{\hat{F}_n(d_{s_1}) > p + \Delta p_2 \mid S_1 = s_1\} = 1. \quad (6)$$

Given the interval design's transition rules, this means that eventually the next-lower level, d_{s_1-1} , will be allocated following each visit to d_{s_1} with a conditional probability of 1. It follows that $(s_1 - 1) \in \mathbb{S}$, reaching a contradiction. We conclude that there is no $s_1 > j^*$ for which $\Pr(S_1 = s_1) > 0$, and therefore one cannot condition on such an event as was done in (6) – and similarly, no $s_2 < j^*$ for which $\Pr(S_2 = s_2) > 0$.

Now we can safely assume that $j^* \in \mathbb{S}$, since this event has probability 1. Given the theorem's conditions and according to Lemma 1, eventually for n large enough

$$\Pr(\hat{F}_n(d_{j^*}) \in (p - \Delta p_1, p + \Delta p_2)) = 1,$$

and so upon the next visit to d_{j^*} it will be repeatedly allocated with probability 1. This means, that with probability one there can be no other level in $\mathbb{S}$. Therefore $\Pr(A) = 1$, and the interval design converges almost surely.

(ii) In the same vein, if all true F values of design levels are above or below the target interval, then with probability 1 the boundary level with F value closest to the interval (d_1 or d_K) belongs to $\mathbb{S}$, and eventually the design will repeatedly hit upon the boundary condition mandating repetition of that level with probability 1. $\square$

References

- Azriel, D., M. Mandel, and Y. Rinott (2011): "The treatment versus experimentation dilemma in dose finding studies," *J. Statist. Plann. Inference*, 141, 2759–2768.
- Babb, J., A. Rogatko, and S. Zacks (1998): "Cancer phase I clinical trials: Efficient dose escalation with overdose control," *Stat. Med.*, 17, 1103–1120.
- Chang, H.-H. and Z. Ying (2009): "Nonlinear sequential designs for logistic item response theory models with applications to computerized adaptive tests," *Ann. Stat.*, 37, 1466–1488.
- Cheung, Y. (2002): "Letter to the editor," *Biometrics*, 58, 237–240.
- Cheung, Y. and R. Chappell (2002): "A simple technique to evaluate model sensitivity in the continual reassessment method," *Biometrics*, 58, 671–674.
- Derman, C. (1957): "Non-parametric up-and-down experimentation," *Ann. Math. Statist.*, 28, 795–798.
- Dixon, W. and A. Mood (1948): "A method for obtaining and analyzing sensitivity data," *J. Am. Statist. Assoc.*, 43, 109–126.
- Durham, S. and N. Flournoy (1995): "Up-and-down designs. I. Stationary treatment distributions," in *Adaptive designs (South Hadley, MA, 1992)*, *IMS Lecture Notes Monogr. Ser.*, volume 25, Hayward, CA: Inst. Math. Statist., 139–157.
- Gezmu, M. (1996): *The Geometric Up-and-Down Design for Allocating Dosage Levels*, Ph.D. thesis, American University, Washington, DC, USA.
- Gezmu, M. and N. Flournoy (2006): "Group up-and-down designs for dose-finding," *J. Statist. Plann. Inference*, 136, 1749–1764.
- Haines, L., I. Perevozskaya, and W. Rosenberger (2003): "Bayesian optimal designs for phase I clinical trials," *Biometrics*, 59, 591–600.
- Ivanova, A., N. Flournoy, and Y. Chung (2007): "Cumulative cohort design for dose-finding," *J. Statist. Plann. Inference*, 137, 2316–2327.
- Leung, D. H.-Y. and Y.-G. Wang (2001): "Isotonic designs for phase I trials," *Cont. Clin. Trials*, 22, 126–138.
- Lord, M. (1980): *Applications of Item Response Theory to Practical Testing Problems*, Erlbaum, Hillsdale, NJ.
- Mathew, P. et al. (2004): "Platelet-derived growth factor receptor inhibitor imatinib mesylate and docetaxel: A modular phase I trial in androgen-independent prostate cancer," *J. Clin. Oncol.*, 22, 3323–3329.
- O'Quigley, J., M. Pepe, and L. Fisher (1990): "Continual reassessment method: a practical design for Phase I clinical trials in cancer," *Biometrics*, 46, 33–48.
- O'Quigley, J. and S. Zohar (2006): "Experimental designs for phase I and phase I/II dose-finding studies," *Brit. J. Canc.*, 94, 609–613.

- Oron, A. (2007): *Up-and-Down and the Percentile-Finding Problem*, Ph.D. thesis, University of Washington, Seattle, Washington, USA, available on <http://arxiv.org/abs/0808.3004>.
- Oron, A. and P. Hoff (2009): “The k-in-a-row up-and-down design, revisited,” *Stat. Med.*, 28, 1805–1820.
- Paoletti, X., J. O’Quigley, and J. Maccario (2004): “Design efficiency in dose finding studies,” *Comp. Stat. Dat. Anal.*, 45, 197–214.
- Pisters, P. W. et al. (2004): “Phase I trial of preoperative doxorubicin-based concurrent chemoradiation and surgical resection for localized extremity and body wall soft tissue sarcomas,” *J Clin Oncol*, 22, 3375–3380.
- R Development Core Team (2011): *R: A language and environment for statistical computing*, R Foundation for Statistical Computing, Vienna, Austria, ISBN 3-900051-07-0.
- Robertson, T., F. T. Wright, and R. L. Dykstra (1988): *Order restricted statistical inference*, Wiley Series in Probability and Mathematical Statistics: Probability and Mathematical Statistics, Chichester: John Wiley & Sons.
- Roy, A., S. Ghosal, and W. Rosenberger (2009): “Convergence properties of sequential Bayesian D-optimal designs,” *J. Statist. Plann. Inference*, 139, 425–440.
- Shen, L. Z. and J. O’Quigley (1996): “Consistency of continual reassessment method under model misspecification,” *Biometrika*, 83, 395–405.
- Shirayev, A. (1996): *Probability, Second Edition*, New York: Springer, translated by R.P. Boas.
- Watson, A. and D. Pelli (1979): “The QUEST staircase procedure,” *Appl. Vis. News.*, 14, 6–7.
- Yuan, Z. and R. Chappell (2004): “Isotonic designs for phase I cancer clinical trials with multiple risk groups,” *Clin. Trials*, 1, 499–508.
- Zacks, S., A. Rogatko, and J. Babb (1998): “Optimal bayesian-feasible dose escalation for cancer Phase I trials,” *Stat. Prob. Lett.*, 38, 215–220.