Using Limited Trial Evidence to Credibly Choose Treatment Dosage when Efficacy and Adverse Effects Weakly Increase with Dose

Charles F Manski
Department of Economics and Institute for Policy Research, Northwestern University

April 2024

Abstract

It has become standard in medical treatment to base dosage on evidence in randomized trials. Yet it has been rare to study how outcomes vary with dosage. In trials to obtain drug approval, the norm has been to compare some dose of a new drug with an established therapy or placebo. Standard trial analysis views each trial arm as qualitatively different, but it may be credible to assume that efficacy and adverse effects (AEs) weakly increase with dosage. Optimization of patient care requires joint attention to both, as well as to treatment cost. This paper develops methodology to use limited trial evidence to choose dosage when efficacy and AEs weakly increase with dose. I suppose that dosage is an integer $t \in (0,1, . ,T)$, T being a specified maximum dose. I study dosage choice when trial evidence on outcomes is available for only K dose levels, where $K < T+1$. Then the population distribution of dose response is partially identified. I show that the identification region is a convex polygon. I characterize clinical and population decision making using the minimax-regret criterion. A simple analytical solution exists when $T = 2$. Computation is tractable when T is larger.

I have benefited from the opportunity to present this work in seminars at Cornell University and Northwestern University. John Mullahy, Mark Ratain, and anonymous reviewers have provided helpful comments.

The need to choose a treatment dosage arises in many medical settings. It has become standard to base dosage on evidence in randomized trials. Yet trial evidence has commonly been limited to comparison of at most a few dose levels.

Evidence on response to dosage has been particularly limited in Phase III trials performed to obtain FDA approval to market new drugs. These trials rarely study how patient outcomes vary with dosage. They typically specify some dose of a new drug and compare it with an established therapy or placebo. When evaluating cancer drugs, the specified dose is often the *maximum tolerated dose*, defined by National Cancer Institute(2023) as "the highest dose with acceptable side effects." To determine the maximum tolerated dose, pharmaceutical firms may compare dose levels in "dose-finding trials" in Phase II studies used to guide design of Phase III trials (Viele *et al.* 2015). However, Phase II sample sizes are usually small and findings unpublished.

Illustration: Clinicians choosing adjuvant care for various cancers may select between immunotherapy and surveillance. Immunotherapy aims to lower the risk of disease recurrence, but it may generate adverse effects (AEs). It is credible that, as dosage increases, the risk of recurrence falls and that of AEs rises. The reason is that stimulation of the immune system may destroy both malignant and healthy cells. Increasing dosage is also more costly. The optimal dosage appropriately weighs these effects, which may be patient specific. The standard practice in Phase III studies has been to perform a two-armed trial evaluating a specified dose of a new immunotherapy. Examples include trials evaluating immunotherapies for melanoma (Eggermont *et al*. 2015, Weber *et al*. 2017, Eggermont et al. 2018). These trials provided no direct information about outcomes with alternative dosages. Modeling dose response suggests that the dose levels specified in the trials could be lowered by at least fifty percent without reduction of efficacy (Ratain and Goldstein 2018; Peer *et al*. 2022). ■

I wrote that trials provide no "direct" information about alternative dosages. The qualifier "direct" stems from the fact that standard trial analysis views each treatment arm as qualitatively different. Consider a trial comparing treatments (A,B,C). In standard analysis, findings with arm A are not used to draw conclusions about B and C, and vice versa. A recent study moving away from this aspect of standard analysis is Stensrud *et al.* (2022).

Standard analysis is well-grounded when treatment arms differ qualitatively. However, this commonly is not so when comparing dosages. If (A,B,C) are increasing dosages, there is often reason to think that efficacy and AEs both increase with dosage. Then findings for arm A provide indirect information about B and C, yielding bounds on outcomes. Standard analysis does not use this information.

A further shortcoming of standard trial analysis has been the practice of studying efficacy and AEs separately, viewing the former as the primary outcome and the latter as secondary outcomes. Optimization of care requires joint attention to efficacy and AEs. The tension between efficacy and AEs makes it impossible to find a dosage that simultaneously maximizes therapeutic effect while minimizing toxicity. The optimal dosage must appropriately weigh these forces.

The long-run solution to the dearth of trial evidence is to perform new trials that enrich the available data. There is scant reason to expect that pharmaceutical firms will voluntarily enhance Phase III trials to have multiple dosage arms. Given a fixed total sample size, a multi-armed trial reduces sample size per arm, lowering statistical power and making it harder to obtain drug approval. Moreover, firms seeking approval of cancer drugs have an incentive to seek approval of the designated maximum tolerated dose rather than a less profitable smaller dose.

With sufficient funding, clinical researchers could perform informative new dosage trials. This has been rare, but there have been some efforts. I describe below a notable set of trials studying dosage of an adjuvant immunotherapy for patients with a particular type of breast cancer.

Adjuvant trastuzumab for HER2-positive early breast cancer: Following several trials comparing trastuzumab treatment for 12 months with no trastuzumab (Piccart-Gebhart *et al*. 2005; Raymond *et al*.

2005; Slamon *et al*. 2011), 12-month treatment duration became standard. However, concern with the cardiotoxicity and cost of trastuzumab generated interest among oncologists in the possibility that a shorter duration may be preferable. Subsequently, multiple two-armed trials comparing 12-month duration with 6 -month, 12-week, or 9-week duration have been performed. Encouraging findings in the FinHer trial (Joensuu *et al*. 2006), which compared 9-week trastuzumab with no trastuzumab, stimulated larger trials comparing other durations. These include the much larger PERSEPHONE trial (Earl *et al*. 2019), comparing 6-month and 12-month trastuzumab. The norm in analysis of these trials has been to measure efficacy primarily by disease-free survival and AEs mainly by cardiotoxocity. Unfortunately, the journal articles reporting on these trials follow the standard practice of studying efficacy and AEs separately rather than jointly. Hence, it is not straightforward to use the reported findings to assess patient welfare. ■

For now, it is important to use existing data effectively. This goal motivates the present paper. I develop methodology to credibly use limited trial evidence to choose treatment dosage when efficacy and AEs weakly increase with dose. The methodology combines partial identification analysis and minimax-regret (MMR) decision making. See Manski(2007) for exposition of both subjects. Some analyses of partial identification in epidemiology include Manski(2021a), Li et al.(2023), and Diemer et al.(2024). Studies of MMR decisions in epidemiology include Manski and Tetenov(2019, 2021).

I suppose that dosage is an integer $t \in (0, . ,T)$, where T is a specified maximum dose. I assume the objective is to maximize mean patient welfare, which is a function of treatment efficacy and AEs, minus cost. To simplify analysis, I consider illness and AEs to be binary events rather than outcomes that vary in severity. I suppose that dosage is a one-time choice rather than a dynamic decision.

The difficulty addressed in the paper is that trial evidence is available for only $K < T+1$ dose levels. Assuming only that efficacy and AEs weakly increase with dosage, the population distribution of dose response is partially identified. I find that the identification region is a convex polygon determined by linear equalities and inequalities.

The identification problem makes it infeasible to choose dosage to maximize patient welfare. Nevertheless, a clinician or a population health planner can make a reasonable choice under ambiguity. A simple MMR solution exists when T = 2 and computation is tractable when T is larger.

This paper differs substantially from traditional research on dose response estimating a parametric model of a univariate outcome. A leading case in pharmacology is nonlinear least squares estimation of a certain three-parameter version of the *Hill equation* predicting a univariate biological response generated by a scalar drug concentration. See Gesztelyi *et al.*(2012), equation 8.

Although monotonicity assumptions have been used in many branches of statistical theory, their use in this paper differs substantially from elsewhere. As far as I am aware, the present analysis is only related to the study of monotone treatment response in Manski(1997). That article analyzed partial identification using observational data when a univariate outcome varies monotonically with treatment intensity. It did not study decision making. The present setting differs most notably because the (efficacy,AE) outcome is bivariate, with increasing dosage improving the former but worsening the latter. Identification analysis with bivariate outcome monotonicity is more complex than with univariate monotonicity.

THE DOSAGE-CHOICE PROBLEM

Let j label a patient and let two dose-dependent treatment outcomes $[d_j(t),e_j(t)]$ determine patient welfare. Given dose t, patient j may experience a disease [$d_j(t) = 1$ if yes, $d_j(t) = 0$ if no] and/or an AE [$e_j(t) = 1$ if yes, $e_j(t) = 0$ if no]. Thus, $[d_j(t),e_j(t)]$ takes one of the values [(0,0), (1,0), (0,1), (1,1)].

Welfare is a function of these outcomes, denoted $w_j[d_j(t),e_j(t)]$. Cost, including monetary and non-monetary costs of administering treatment, is a function $g_j(t)$ of dose level, measured in the same units as welfare. I index outcomes, welfare, and cost by j, permitting heterogeneity across the population.

Assume that the objective is to maximize welfare net of cost. Interpretation of optimal dosage depends on the available knowledge. An idealized setting assumes perfect foresight, knowing $w_j[d_j(t),e_j(t)]$ and $g_j(t)$ for $t \in \{0, . ,T\}$. Then an optimal dose for patient j solves $\max_{t = 0, . , T} w_j[d_j(t),e_j(t)] - g_j(t)$.

Perfect foresight is too unrealistic to provide a useful benchmark for decision making. It has been common in medical economics to study optimal utilitarian clinical care assuming objectively correct probabilistic expectations; e.g., Phelps and Mushlin(1988). Assume that a clinician observes patient covariates such as age, sex, and health history. Viewing j as a member of a population J of patients who have the same covariates, assume that the clinician knows the dose-dependent distributions p{w[d(t),e(t)], g(t)}, $t \in \{0, . ,T\}$ of (welfare,cost) over J. Then an optimal dose solves $\max_{t=0,..,T} E\{w[d(t),e(t)]\} - E[g(t)]$.

The above knowledge scenario is still unrealistic. Standard data collection in trials yields evidence on disease and AEs, but not on welfare as a function of these outcomes. Medical economists sometimes conduct separate studies that sample patients and aim to learn their welfare functions, by questioning them about the choices they would make if they were to experience hypothetical disease and AEs; e.g., Basu and Meltzer(2007) and Devlin and Brooks(2017). However, the patients in these studies typically are different than the subjects in trials.

Estimation of dose-dependent mean welfare may be feasible if it is credible to assume that welfare functions are mean-independent of disease and AE outcomes; that is, E[w(h, i)|d(t) = h, e(t) = i] = E[w(h,i)], h = 0,1 and i = 0,1. Then mean welfare with dose t is

$$(1)\quad E\{w[d(t),e(t)]\} = E[w(0,0)]\cdot p[d(t)=0, e(t)=0] + E[w(1,0)]\cdot p[d(t)=1, e(t)=0] + E[w(0,1)]\cdot p[d(t)=0, e(t)=1] + E[w(1,1)]\cdot p[d(t)=1, e(t)=1] - E[g(t)].$$

A study measuring patient welfare in hypothetical scenarios can enable estimation of the outcome-specific mean welfare function E[w(·,·)] = {E[w(1,1)], E[w(1,0)], E[w(0,1)], E[w(0,0)]}. A (T+1)-armed trial enables estimation of the dose-dependent outcome distributions p[d(t),e(t)], $t \in \{0, . ,T\}$.

I henceforth assume that (1) holds and that research reveals E[w(·,·)]. I also assume that mean treatment cost E[g(t)] is known. My concern is that it has been rare to perform the aforementioned (T+1)-armed trial. Instead, for some K with $1 < K < T+1$, the practice has been to perform a K-armed trial, assigning subjects

to dosages t_k, k = 1, . , K. With this design, a trial directly enables one to estimate the outcome distributions $p[d(t_k),e(t_k)]$, k = 1, . , K, but not p[d(t),e(t)] for other dosages. Supposing that the trial is large enough that sampling imprecision is negligible, standard trial analysis enables solution of the dose-constrained optimization problem $\max_{k = 1, . , K} E\{w[d(t_k),e(t_k)]\} - E[g(t_k)]$, but not the unconstrained problem $\max_{t = 0, . , T} E\{w[d(t),e(t)]\} - E[g(t)]$.

IDENTIFICATION WITH MONOTONE DOSE RESPONSE

Monotone Dose Response

It is often credible that efficacy and AEs increase with dosage. Thus, a patient who is disease-free with a low dose would remain disease-free with a higher dose. A patient who has an AE with a low dose would have one with a higher dose. Formally, I assume monotone dose response:

Monotone Dose Response: Consider doses (s,t) with t > s. For each patient j,

Monotone Efficacy (ME): $d_j(s) = 0 \Rightarrow d_j(t) = 0$.

Monotone AEs (MT): $e_j(s) = 1 \Rightarrow e_j(t) = 1$.

An alternative way to express monotone dose response uses the concept of patient-specific threshold dose levels, as follows:

Monotone Dose Response Expressed Through Threshold Dose Levels: For each patient j,

Monotone Efficacy (ME): Either $d_j(t) = 1$, all $t \in \{0, . , T\}$, or there exists a threshold dose $t_{dj} \in \{0, . ,T\}$ such that $d_j(t) = 1$ for $t < t_{dj}$ and $d_j(t) = 0$ for $t \geq t_{dj}$. When $d_j(t) = 1$, all $t \in \{0, . ,T\}$, it suffices to define an infeasible threshold dose $t_{dj} = T+1$.

Monotone AEs (MT): Either $e_j(t) = 0$, all $t \in \{0, . , T\}$ or there exists a threshold dose $t_{ej} \in \{0, . ,T\}$ such that $e_j(t) = 0$ for $t < t_{ej}$ and $e_j(t) = 1$ for $t \geq t_{ej}$. When $e_j(t) = 0$, all $t \in \{0, . ,T\}$, it suffices to define an infeasible threshold dose $e_{dj} = T+1$.

The two versions of (ME)−(MT) are equivalent, as $d_j(t) = 1[t < t_{dj}]$ and $e_j(t) = 1[t \geq t_{ej}]$.

While assumptions ME and MT are often credible, it would go too far to assert that they hold universally. Either may not hold if there exist biological interactions between disease and AEs. Disease may weaken a patient, increasing susceptibility to AEs. AEs may weaken a patient, increasing susceptibility to disease. Biological interactions do not necessarily falsify the assumptions. However, they call for careful scrutiny by clinicians knowledgeable about the processes generating disease and AEs.

Identification

Let $q(t_d,t_e)$ denote the population distribution of threshold dose levels. Without trial evidence, $q(t_d,t_e)$ may be any bivariate distribution on $\{0, . ,T+1\}$ x $\{0, . ,T+1\}$. Thus, there are $(T+2)^2$ component probabilities. Proposition 1 gives the findings with trial evidence. I denote the identification region for $q(t_d,t_e)$ as Q. The findings for identification of outcome distributions and mean welfare follow from determination of Q. The identification regions for $p[d(t),e(t)]$, $t \in \{0, . ,T\}$ and $E\{w[d(t),e(t)]\}$, $t \in \{0, . ,T\}$ are denoted P and W.

Proposition 1: Let (ME) and (MT) hold. Let a K-armed trial assigning subjects to dosages t_k, $k = 1, . ,K$ reveal $p[d(t_k),e(t_k)]$, $k = 1, . ,K$. Then Q, P, and W are the sets of distributions derived in eAppendix 1. ■

Proposition 1 shows that $q(t_d,t_e)$ is partially identified. Recall that $q(t_d,t_e)$ has $(T+2)^2$ component probabilities. eAppendix 1 shows that these probabilities solve 4K+1 linear equations. Of these equations, 3K+1 are linearly independent. Hence, $q(t_d,t_e)$ solves 3K+1 non-redundant linear equations in $(T+2)^2$

unknowns, plus the non-negativity inequalities required of all probabilities. Thus, Q is a $(T+2)^2 - (3K+1)$ dimensional convex polygon.

eAppendix 1 shows that dose-dependent outcome distributions are linear functions of $q(t_d,t_e)$. Hence, P is a convex polygon. Mean welfare is a linear function of the probabilities p. Hence, W is an interval on the real line. Computation of the identification regions is straightforward. Q is a convex polygon determined by linear equalities and inequalities. P and W are linear transformations of Q.

This analysis of identification assumes only monotone dose response, without restricting the distribution $q(t_d,t_e)$. Assumptions on this distribution may have identifying power and be credible in some applications. eAppendix 2 discusses some possibilities.

DOSAGE CHOICE WITH MONOTONE DOSE RESPONSE

Consider dosage choice when a trial with dosages t_k, k = 1, . ,K reveal $p[d(t_k),e(t_k)]$, k = 1, . ,K. I consider a clinician treating one patient and a health planner treating a population. The clinician must choose a single treatment for the patient and cannot randomize the choice. The planner can choose a fractional treatment allocation $\delta(t)$, $t \in \{0, . ,T\}$ in the unit simplex Δ on R^{T+1}. Thus, the planner has a richer set of options than the clinician.

If a utilitarian clinician and planner were to have objectively correct probabilistic expectations, they would make the same optimal dosage choice for each patient. However, they may behave differently with partial knowledge of $q(t_d,t_e)$, which makes optimal choice infeasible. Being able to make a fractional treatment allocation, a planner can diversify treatment, which limits the magnitude of treatment errors. A clinician treating one patient cannot diversify. Manski (2009) provides discussion and analysis.

The Bayesian prescription is to assert a subjective probability distribution on the $(T+2)^2$ -dimensional threshold-dose distribution $q(t_d,t_e)$ and maximize subjective expected welfare. Objective mean welfare $E\{w[d(t),e(t)]\}$ is linear in $q(t_d,t_e)$. A Bayesian asserts a subjective distribution, say π, on the space of all possible threshold dose distributions. If the trial sample size is large enough to ignore sampling imprecision,

the space of possible $q(t_d,t_e)$ given assumptions ME and MT is the identification region Q. Thus, a Bayesian places a subjective distribution on $q(t_d,t_e)$ posterior to performance of identification analysis.

Using the subjective distribution, a Bayesian replaces the unknown $q(t_d,t_e)$ with its subjective expectation $\int q(t_d,t_e)d\pi$, and chooses the dose that would be optimal if $\int q(t_d,t_e)d\pi$ were the actual threshold dose distribution. Bayesian decision making is simple, but it is sensible only if the decision maker places a credible subjective distribution on $q(t_d,t_e)$. Bayesians have long struggled to provide guidance on credible specification of subjective distributions. The matter continues to be controversial; see Spiegelhalter *et al.*(1994).

Inability to specify a credible subjective distribution motivates study of decision making under ambiguity, which does not place a subjective distribution on $q(t_d,t_e)$. Two prominent approaches are maximin and minimax regret (MMR). I evaluate options by their maximum regret.

I study the MMR criterion because, as discussed in Manski(2018,2019,2021b) inter alia, this criterion is conceptually appealing for choice under ambiguity. Regret quantifies how lack of knowledge of the true distribution $q(t_d,t_e)$ diminishes the quality of decisions. The term "maximum regret" is shorthand for the maximum sub-optimality of a decision across feasible states of nature. A decision with small maximum regret is uniformly near-optimal across all states. This is a desirable property.

Minimax-Regret Decisions

I first pose the MMR criterion abstractly and then apply it to dosage choice. Consider a planner who faces choice set C and believes that the true state of nature lies in specified state space S. An objective function $f(\cdot,\cdot)$: $C \times S \to R^1$ maps actions and states into welfare. The planner wants to maximize true welfare but does not know the true state. In settings without statistical imprecision, the MMR criterion solves

$$(2) \qquad \min_{c\in C} \max_{s\in S} [\max_{d\in C} f(d,s) - f(c,s)].$$

$\max_{d \in C} w(d,s) - w(c,s)$ is the *regret* of action c in state s; that is, the degree of suboptimality.

In the present setting, the state space indexes the feasible distributions $q(t_d,t_e)$ derived in Proposition 1; thus, S = Q. In the clinical case, the choice set comprises all feasible doses; thus, C = {0, . , T}. In the planning case, the choice set comprises all fractional allocations; thus, C = Δ.

In the clinical case, the objective function in state s is $f[t, q_s(t_d,t_e)] = E_s\{w[d(t),e(t)]\}$, where $q_s(t_d,t_e)$ is the distribution indexed by state s and E_s denotes expectation with respect to this distribution. In the planning case, the objective function is the expectation of the mean welfare function with respect to the treatment allocation; thus, $f[\delta(\cdot), q_s(t_d,t_e)] = \sum_{t=0..,T} \delta(t)\cdot E_s\{w[d(t),e(t)]\}$.

Computation

Exact computation of the MMR decision is tractable in the clinical case when T is not extremely large. An alternative expression of (2) reverses the two max operations to obtain

$$(2') \qquad \min_{c\in C} \max_{d\in C} \max_{s\in S} [f(d,s) - f(c,s)].$$

Holding c and d fixed, consider the inner maximization over S. When c = d, the maximum is 0. When c ≠ d, f(d,s) − f(c,s) is linear in $q_s(\cdot,\cdot)$ in the setting of this paper. $\max_{s\in S} [f(d,s) - f(c,s)]$ is a linear programming problem, solvable with standard algorithms. Considering all c ≠ d, determination of the MMR dosage requires solution of T·(T+1) linear programming problems.

In the planning case, $\max_{s\in S} [f(d,s) - f(c,s)]$ remains a linear programming problem, but f(d,s) − f(c,s) is a different linear function of $q_s(\cdot,\cdot)$. The set of feasible dosage allocations is the entire simplex Δ on R^{T+1} rather than only the vertices, which place probability one on single treatments. When T > 2, this makes exact computation of the MMR dosage allocation infeasible. However, it is feasible to use a finite grid to approximate Δ and solve the associated finite number of linear programming problems. I am not aware of any sophisticated numerical approach to maximize and minimize the objective function over the simplex Δ; hence the resort to grid search.

Population-Health Dosage Allocation when T = 2

Computation of the planner's MMR allocation is simple when T = 2, K = 2, $t_1 = 0$, and $t_2 = 2$. Thus, there is no trial evidence for t = 1. This case is important in dosage choice for cancer drugs. As discussed in the opening section, Phase III trials commonly compare zero dose and a maximum tolerable dose, without evaluating smaller positive doses.

Welfare net of treatment cost is point identified for t = 0 and t = 2. Let them be denoted $\omega_0 \equiv E\{w[d(0),e(0)]\} - E[g(0)]$ and $\omega_2 \equiv E\{w[d(2),e(2)]\} - E[g(2)]$. Welfare net of treatment cost is partially identified for t = 1, being known to lie in a computable interval, say $[\omega_{1L},\omega_{1U}]$, determined by the identification analysis. I consider settings with $\omega_{1L} < \max(\omega_0,\omega_2) < \omega_{1U}$, where there is ambiguity.

Manski(2009) studied MMR allocation of a population between two undominated treatments. The analysis applies here, where t = 1 and either t = 0 or t = 2 are undominated. The MMR allocation assigns positive fractions of patients to t = 1 and to the other dose with the higher value of ω. If $\omega_2 > \omega_0$, t = 0 is dominated and receives zero allocation. Analysis shows that the fraction of the population assigned to t = 1 is $(\omega_{1U}-\omega_2)/(\omega_{1U}-\omega_{1L})$ and the fraction assigned to t = 2 is $(\omega_2-\omega_{1L})/(\omega_{1U}-\omega_{1L})$. Symmetrically, if $\omega_2 < \omega_0$, t = 2 is dominated and receives zero allocation. The fraction assigned to t = 1 is $(\omega_{1U}-\omega_0)/(\omega_{1U}-\omega_{1L})$ and that assigned to t = 0 is $(\omega_0-\omega_{1L})/(\omega_{1U}-\omega_{1L})$. eAppendix 3 illustrates numerically.

DISCUSSION

This paper has studied dosage choice with limited trial data. The theme is that standard analysis of trial data is unnecessarily restrictive when credible assumptions make findings informative across treatments. The assumption of monotone dose response is often credible in dosage trials.

To apply the analysis to existing trial data, it is necessary to observe the joint empirical distribution of (efficacy,AE) outcomes. Unfortunately, the standard medical research practice has been to report findings on efficacy and AEs separately. The joint outcome distribution would be computable if trial investigators were to make individual patient data available to other researchers. For example, the existing trial data

comparing trastuzumab doses could be analyzed from the perspective of this paper if individual patient data were available.

It would be valuable to extend the analysis in multiple ways. eAppendix 4 discusses MMR decision making with sample data. Another direction is to recognize that disease and AE outcomes may vary in severity rather than being binary. When outcomes vary in severity, the distribution of outcomes may still be characterized by a threshold dose distribution $q\{t_d.t_e)$, but now t_d and t_e is each a vector of thresholds for passage into increasingly high levels of severity.

Another direction is to recognize that the scalar dosing decision studied here is often a component of a more general problem of choice of a (dose, schedule, duration) regimen for administration of treatment. The concept of monotone treatment extends to regimens. One regimen is more intense than another if it combines a weakly larger dose, more frequent schedule, and longer duration. Analysis of regimens is complex when comparing alternatives that are more intense in some components but less in others.

Yet another direction is to consider multiperiod treatment settings, where a planner chooses dosage for a sequence of patient cohorts. Learning is then possible, with observation of the outcomes experienced by earlier cohorts informing dosage choice for later cohorts. Fractional dosage allocations are advantageous for learning because they generate randomized experiments yielding outcome data on multiple dosages. This suggests *adaptive diversification* of dosage, discussed in general treatment contexts in Manski(2009).

References

Basu, A., and D. Meltzer (2007), "Value of Information on Preference Heterogeneity and Individualized Care," *Medical Decision Making*, 27, 112–127.

Devlin, N. and R. Brooks (2017), "EQ-5D and the EuroQol Group: Past, Present and Future," *Applied Health Economics and Health Policy*, 15, 127-137.

Diemer, E., J. Shi, and S. Swanson (2024), "Partial Identification of the Effects of Sustained Treatment Strategies," *Epidemiology*, DOI: 10.1097/EDE.0000000000001721 .

Earl, H., L.Hiller, A. Vallier *et al*. (2019), "6 versus 12 months of Adjuvant Trastuzumab for HER2-Positive Early Breast Cancer (PERSEPHONE): 4-year disease-free survival results of a randomised phase 3 non-inferiority trial," *Lancet*, 393, 2599-2612.

Eggermont A., C. Blank, M. Mandala *et al.* (2018), "Adjuvant Pembrolizumab versus Placebo in Resected Stage III Melanoma," *New England Journal of Medicine,* 378, 1789-1801.

Eggermont A., V. Chiarion-Sileni, J. Grob J *et al.* (2015), "Adjuvant Ipilimumab versus Placebo after Complete Resection of High-Risk Stage III Melanoma (EORTC 18071): a randomised, double-blind, phase 3 trial," *Lancet Oncology*, 16, 522-30.

Gesztelyi, R., J. Zsuga, A. Kemeny-Beke, B. Varga. B. Juhasz, and A. Tosaki (2012), "The Hill Equation and the Origin of Quantitative Pharmacology," *Archive for History of Exact Sciences*, 66, 427–438.

Joensuu H., P. Kellokumpu-Lehtinen, P. Bono *et al*. (2006), "Adjuvant Docetaxel or Vinorelbine with or without Trastuzumab for Breast Cancer," *New England Journal of Medicine,* 356, 809-820.

Li, S., V. Litvin, and C. Manski (2023), "Partial Identification of Personalized Treatment Response with Trial-reported Analyses of Binary Subgroups," *Epidemiology*, 34, 319-324.

Manski C. (1997), "Monotone Treatment Response," *Econometrica*, 65, 1311-1334.

Manski, C. (2007), *Identification for Prediction and Decision*, Cambridge: Harvard University Press.

Manski, C. (2009), "Diversified Treatment under Ambiguity," *International Economic Review*, 50, 1013-1041.

Manski, C. (2018), "Reasonable Patient Care under Uncertainty," *Health Economics*, 27, 1397-1421.

Manski, C. (2019), *Patient Care under Uncertainty*, Princeton, NJ: Princeton University Press.

Manski, C. (2021a), "Bounding the Accuracy of Diagnostic Tests, with Application to COVID-19 Antibody Tests," *Epidemiology*, 32, 162-167.

Manski, C. (2021b), "Econometrics for Decision Making: Building Foundations Sketched by Haavelmo and Wald," *Econometrica* , 89, 2827-2853.

Manski, C. and A. Tetenov (2019), "Trial Size for Near-Optimal Choice Between Surveillance and Aggressive Treatment: Reconsidering MSLTII," *The American Statistician*, 73, 305-311.

Manski, C. and A. Tetenov (2021), "Statistical Decision Properties of Imprecise Trials Assessing Coronavirus 2019 (COVID-19) Drugs," *Value in Health*, 24, 641-647.

National Cancer Institute (2023), *NCI Dictionary of Cancer Terms*, https://www.cancer.gov/publications/dictionaries/cancer-terms, accessed April 23, 2023.

Peer, C., B. Heiss, D. Goldstein, J. Goodell, W. Figg, and M. Ratain (2022), "Pharmacokinetic Simulation Analysis of Less Frequent Nivolumab and Pembrolizumab Dosing: Pharmacoeconomic Rationale for Dose Deescalation," *Journal of Clinical Pharmacology*, 62, 532-540.

Phelps, C., and A. Mushlin (1988), "Focusing Technology Assessment Using Medical Decision Theory." *Medical Decision Making*, 8, 279-289.

Piccart-Gebhart M., M. Procter, B. Leyland-Jones *et al*. (2005), "Trastuzumab after Adjuvant Chemotherapy in HER2-Positive Breast Cancer," *New England Journal of Medicine,* 353*,* 1659-1672.

Ratain M. and D. Goldstein (2018), "Time is Money: Optimizing the Scheduling of Nivolumab," *Journal of Clinical Oncology*, 36, 3074-3076.

Romond, E., E. Perez, J. Bryant *et al*. (2005), "Trastuzumab plus Adjuvant Chemotherapy for Operable HER2-Positive Breast Cancer," *New England Journal of Medicine***,** 353**,** 1673-1684.

Slamon D., W. Eiermann, N. Robert *et al*. (2011), "Adjuvant Trastuzumab in HER2-Positive Breast Cancer," *New England Journal of Medicine,* 365, 1273-1283.

Spiegelhalter D., L.Freedman, and M. Parmar (1994), "Bayesian Approaches to Randomized Trials (with discussion). *Journal of the Royal Statistics Society Series A*, 157, 357-416.

Stensrud, M., J. Young, V. Didelez, J. Robins, and M. Hernán (2022). "Separable Effects for Causal Inference in the Presence of Competing Events," *Journal of the American Statistical Association*, 117, 175-183.

Viele K. and J. Connor (2015), "Dose-Finding Trials Optimizing Phase 2 Data in the Drug Development Process," *Journal of the American Medical Association*, 314, 2294-2295.

Weber J., M. Mandala, M. Del Vecchio *et al.* (2017), "Adjuvant Nivolumab versus Ipilimumab in Resected Stage III or IV Melanoma," *New England Journal of Medicine*, 377, 1824-1835.

eAPPENDIX 1. PROOF OF PROPOSITION 1

The non-negativity of probabilities and the Law of Total Probability imply that $q(t_d, t_e)$ satisfies these inequalities and equality:

(A1) $q(t_d = h, t_e = i) \geq 0, h \in \{0, . , T + 1\}, i \in \{0, . , T + 1\}.$

(A2) $1 = \sum_{h=0}^{T+1} \sum_{i=0}^{T+1} q(t_d = h, t_e = i).$

The dose-dependent outcome distributions $p[d(t), e(t)]$, $t \in \{0, . , T\}$ have $4(T + 1)$ component probabilities. Equations (A3a)−(A3d) express outcome probabilities in terms of $q(t_d, t_e)$. For $t \in \{0, . , T\}$,

(A3a) $p[d(t) = 0, e(t) = 0] = q(t_d \leq t, t_e > t) = \sum_{h=0}^{t} \sum_{i=t+1}^{T+1} q(t_d = h, t_e = i).$

(A3b) $p[d(t) = 1, e(t) = 0] = q(t_d > t, t_e > t) = \sum_{h=t+1}^{T+1} \sum_{i=t+1}^{T+1} q(t_d = h, t_e = i).$

(A3c) $p[d(t) = 0, e(t) = 1] = q(t_d \leq t, t_e \leq t) = \sum_{h=0}^{t} \sum_{i=0}^{t} q(t_d = h, t_e = i).$

(A3d) $p[d(t) = 1, e(t) = 1] = q(t_d > t, t_e \leq t) = \sum_{h=t+1}^{T+1} \sum_{i=0}^{t} q(t_d = h, t_e = i).$

These preliminaries yield the identification regions Q, P, and W. Q comprises all distributions $q(t_d, t_e)$ that satisfy (A1), (A2), and, for $k = 1, . , K$,

(A4a) $p[d(t_k) = 0, e(t_k) = 0] = q(t_d \leq t_k, t_e > t_k) = \sum_{h=0}^{t_k} \sum_{i=t_k+1}^{T+1} q(t_d = h, t_e = i),$

(A4b) $p[d(t_k) = 1, e(t_k) = 0] = q(t_d > t_k, t_e > t_k) = \sum_{h=t_k+1}^{T+1} \sum_{i=t_k+1}^{T+1} q(t_d = h, t_e = i),$

$$\text{(A4c)}\quad p[d(t_k) = 0, e(t_k) = 1] = q(t_d \le t_k, t_e \le t_k) = \sum_{h=0}^{t_k} \sum_{i=0}^{t_k} q(t_d = h, t_e = i),$$

$$\text{(A4d)}\quad p[d(t_k) = 1, e(t_k) = 1] = q(t_d > t_k, t_e \le t_k) = \sum_{h=t_k+1}^{T+1} \sum_{i=0}^{t_k} q(t_d = h, t_e = i).$$

P comprises all distributions p[d(t), e(t)], $t \in \{0, . , T\}$ that satisfy (A3a)−(A3d) for some element of Q. W comprises all means $E\{w[d(t), e(t)]\}$, $t \in \{0, . , T\}$ that satisfy (1) for some element of P.

Q. E. D.

eAPPENDIX 2. RESTRICTING THE DISTRIBUTION OF THRESHOLD DOSES

No AEs Without Treatment: A credible assumption is that AEs cannot occur when the dose level is zero. Formally, assume that $q(t_d = h, t_e = 0) = 0$, $h \in \{0, . , T + 1\}$. The assumption fixes $T + 2$ components of $q(t_d, t_e)$, reducing the number of unknown components from $(T + 2)^2$ to $(T + 2)(T + 1)$. It implies that $p[d(0) = 0, e(0) = 1] = p[d(0) = 1, e(0) = 1] = 0$.

Concurrent Thresholds for Efficacy and AEs: In some contexts, it is credible to pose assumptions on the cross-patient association between thresholds t_d and t_e. A positive association is credible in the treatment of cancer with immunotherapy because stimulation of the immune system increases with dose level, making both tumor reduction and AEs more likely. Patient immune systems vary genetically and epigenetically, so the degree to which immunotherapy generates immune responses may vary across patients.

There are many ways to formalize positive association as an assumption on $q(t_d, t_e)$. A polar case assumes that each patient has one threshold dose that jointly prevents disease recurrence and generates an AE. That is, there exists a patient-specific dose t_{*j} such that $t_{*j} = t_{dj} = t_{ej}$, all j. Then $q(t_d = h, t_e = i) = 0$

whenever $h \neq i$, reducing the number of unknown components of $q(t_d, t_e)$ from $(T + 2)^2$ to $T + 2$. The assumption implies that, for each dose level t, $p[d(t) = 0, e(t) = 0] = p[d(t) = 1, e(t) = 1] = 0$.

Statistical Independence of Efficacy and AE Thresholds: Whereas the biological mechanism of immunotherapy suggests positive association between thresholds, in other treatment contexts separate physiological processes may make a treatment efficacious and yield AEs. In these cases, it may be reasonable to assume statistical independence of t_d and t_e. Then $q(t_d = h, t_e = i) = q(t_d = h) \cdot q(t_e = i)$ for all (h, i), reducing the number of unknown components of $q(t_d, t_e)$ from $(T + 2)^2$ to $2(T + 2)$. The assumption implies that, for each dose level t, $p[d(t), e(t)] = p[d(t)] \cdot p[e(t)]$.

eAPPENDIX 3. NUMERICAL ILLUSTRATIONS

Let mean patient welfare with each (d, e) outcome be as follows:

$$E[w(0, 0)] = 1, E[w(1, 0)] = 0.25, E[w(0, 1)] = 0.75, E[w(1, 1)] = 0.$$

This quantifies the usual situation where it is best to experience neither disease nor an adverse event, and worst to experience both. Experiencing one or the other is intermediate, disease being more harmful on average than an adverse event.

Let the threshold dose distribution be as follows:

$$q(t_d = h, t_e = 0) = 0, h \in \{0, 1, 2, 3\}; \quad q(t_d = h, t_e = i) = 1/12, h \in \{0, 1, 2, 3\}, i \in \{1, 2, 3\}.$$

The first part of the specification expresses the credible assumption that AEs cannot occur without treatment. The uniform distribution in the second part of the specification is only intended to be illustrative, as the actual distribution will be context specific.

The outcome probabilities implied by this threshold dose distribution are

$$p[d(0) = 0, e(0) = 0] = 0.25 \qquad p[d(0) = 0, e(0) = 1] = 0$$
$$p[d(0) = 1, e(0) = 0] = 0.75 \qquad p[d(0) = 1, e(0) = 1] = 0.$$

$$p[d(1) = 0, e(1) = 0] = 0.333 \qquad p[d(1) = 0, e(1) = 1] = 0.167$$

p[d(1) = 1, e(1) = 0] = 0.333 p[d(1) = 1, e(1) = 1] = 0.167.

p[d(2) = 0, e(2) = 0] = 0.25 p[d(2) = 0, e(2) = 1] = 0.5
p[d(2) = 1, e(2) = 0] = 0.083 p[d(2) = 1, e(2) = 1] = 0.166.

The values of mean welfare unconditional on (d, e) are

E{w[d(0), e(0)]} = 0.4375, E{w[d(1), e(1)]} = 0.542, E{w[d(2), e(2)]} = 0.6458.

The values for p[d(0), e(0)], p[d(2), e(2)], E{w[d(0), e(0)]}, and E{w[d(2), e(2)]} are point-identified by the trial evidence. Application of Proposition 1 shows that the identification region for p[d(1), e(1)] is a convex polygon in R^4, whose projections on the four axes are these bounds:

p[d(1) = 0, e(1) = 0] ∈ [0 ,0.75] p[d(1) = 0, e(1) = 1] ∈ [0, 0.5]
p[d(1) = 1, e(1) = 0] ∈ [0.083, 0.75] p[d(1) = 1, e(1) = 1] ∈ [0, 0.67].

The identification region for E{w[d(1), e(1)]} is the interval [0.2708, 0.8125].

Now consider MMR clinical and public-health dosage choice with various cost functions g(t). Here are some illustrative findings:

g(t) = 0, all t: The MMR clinical dose is t = 2, with MMR value 0.167. The public-health dose allocation is (0, 0.308, 0.692) to doses (0, 1, 2), with MMR value 0.116.

g(t) = (0.05)t: The MMR clinical dose is t = 2, with MMR value 0.217. The public-health dose allocation is (0, 0.4, 0.6) to doses (0, 1, 2), with MMR value 0.13.

g(t) = (0.1)t: The MMR clinical dose is t = 2, with MMR value 0.267. The public-health dose allocation is (0, 0.49, 0.51) to doses (0, 1, 2), with MMR value 0.136.

g(t) = (0.15)t: The MMR clinical dose is t = 0, with MMR value 0.225. The public-health dose allocation is (0.59, 0.41, 0) to doses (0, 1, 2), with MMR value 0.132.

g(0) = g(1) = 0, g(2) = 0.30: The MMR clinical dose is t = 1, with MMR value 0.167. The public-health dose allocation is (0.308, 0.692, 0) to doses (0, 1, 2), with MMR value 0.115.

Observe that the clinical decision is t = 2 or t = 0 with each of the four linear cost functions, but it is t = 1 with the nonlinear cost function. Thus, a clinician using the MMR criterion may choose a dose level excluded from the trial design.

eAPPENDIX 4. DECISIONS WITH FINITE-SAMPLE TRIAL DATA

The Wald (1950) framework of statistical decision theory supposes that a decision maker observes data generated by a sampling distribution, which varies with the state of nature. To express this, let Ψ denote the sample space; that is, the set of samples that may be drawn. Let the possible sampling distributions be denoted $(Q_s, s \in S)$. The literature assumes that the sample space does not vary with s and is known, whereas the sampling distribution varies with s and is unknown. A statistical decision function, $c(\cdot): \Psi \rightarrow C$ maps the sample data into a chosen action.

An SDF is a deterministic function after realization of the sample data, but it is a random function ex ante. Hence, an SDF generically makes a randomized choice of an action. Statistical decision theory evaluates the performance of SDF $c(\cdot)$ in state s by $Q_s\{f[c(\psi), s]\}$, the ex-ante distribution of welfare across realizations ψ of the sampling process. In abstraction, the statistical version of the MMR criterion is

$$\min_{c(\cdot) \in \Gamma} \max_{s \in S} (\max_{d \in C} f(d, s) - E_s\{f[c(\psi), s]\}).$$

In the present setting, the choice set and the objective function are as in the main text. Sample data do not reveal the outcome distributions $p[d(t_k), e(t_k)]$, $k = 1, . , K$; the data only enable estimation. Hence, the present state space is larger than in the main text, now comprising all multinomial distributions on $\{0, . , T + 1\} \times \{0, . , T + 1\}$. The trial design draws N(k) persons at random from J and assigns each sampled person to dose level t_k. For $j \in N(k)$, one observes $[d_j(t_k), e_j(t_k)]$. Thus, $\psi = \{N(k), p_{N(k)}[d(t_k), e(t_k)], k = 1, . , K\}$.

Computation of finite-sample MMR decisions is typically a complex problem. I expect this to be so in the setting of this paper. I do not address the problem here, but it is an important topic for future research. A computationally simple approach to dosage choice with sample data is to estimate $p[d(t_k), e(t_k)]$, $k = 1, . , K$ by their sample analogs $p_{N(k)}[d(t_k), e(t_k)]$, $k = 1, . , K$ and then proceed as in the main text, acting as if the estimates are accurate. Manski (2021) discusses this "as-if" approach to decision making in abstraction and explains how to compute its maximum regret numerically.

An interesting question is to learn how the assumption of monotone dose response affects finite-sample inference and decision making, relative to what occurs with standard trial analysis. Even when a complete dosing trial has been performed, enabling standard analysis to point-identify dose response, the monotonicity assumption may still improve finite-sample analysis. It may improve the statistical precision of point estimates of outcome probabilities used in as-if decision making and thereby reduce the maximum regret achieved when estimates are used to choose dosage.

eAppendix References

Manski, C. (2021), "Econometrics for Decision Making: Building Foundations Sketched by Haavelmo and Wald," *Econometrica* , 89, 2827-2853.

Wald, A. (1950), *Statistical Decision Functions,* New York: Wiley.