Costing 'the' MTD

David C. Norris

Precision Methodologies, LLC, Seattle WA, USA

Corresponding author:

David C. Norris

Email address: david@precisionmethods.guru

ABSTRACT

Background: Absent adaptive, individualized dose-finding in early-phase oncology trials, subsequent registration trials risk suboptimal dosing that compromises statistical power and lowers the probability of technical success (PTS) for the investigational drug. While much methodological progress has been made toward adaptive dose-finding, and quantitative modeling of dose-response relationships, most such work continues to be organized around a concept of 'the' maximum tolerated dose (MTD). But a new methodology, Dose Titration Algorithm Tuning (DTAT), now holds forth the promise of *individualized* 'MTD_i' dosing. Relative to such individualized dosing, current 'one-size-fits-all' dosing practices amount to a *constraint* that imposes *costs* on society. This paper estimates the magnitude of these costs.
Methods: Simulated dose titration as in (Norris 2017) is extended to 1000 subjects, yielding an empirical MTD_i distribution to which a gamma density is fitted. Individual-level efficacy, in terms of the probability of achieving remission, is assumed to be an $\text{E}_{\max}$-type function of *dose relative to* MTD_i, scaled (arbitrarily) to identify MTD_i with the LD_{50} of the individual's tumor. (Thus, a criterion 50% of the population achieve remission under individualized dosing in this analysis.) Current practice is modeled such that all patients receive a first-cycle dose at 'the' MTD, and those for whom $\text{MTD}_i < \text{MTD}_{\text{the}}$ experience a 'dose-limiting toxicity' (DLT) that aborts subsequent cycles. Therapy thus terminated is assumed to confer no benefit. Individuals for whom $\text{MTD}_i \geq \text{MTD}_{\text{the}}$ tolerate a full treatment course, and achieve remission with probability determined by the $\text{E}_{\max}$ curve evaluated at $\text{MTD}_{\text{the}}/\text{MTD}_i$. A closed-form expression is obtained for the population remission rate, and maximized numerically over MTD_{the} as a free parameter, thus identifying the best result achievable under one-size-fits-all dosing. A sensitivity analysis is performed, using both a perturbation of the assumed $\text{E}_{\max}$ function, and an antipodal alternative specification.
Results: Simulated MTD_i follow a gamma distribution with shape parameter $\alpha \approx 1.75$. The population remission rate under one-size-fits-all dosing at the maximizing value of MTD_{the} proves to be a function of the shape parameter—and thus the coefficient of variation (CV)—of the gamma distribution of MTD_i. Within a plausible range of $\text{CV}(\text{MTD}_i)$, one-size-fits-all dosing wastes approximately half of the drug's population-level efficacy. In the sensitivity analysis, sensitivity to the perturbation proves to be of second order. The alternative exposure-efficacy specification likewise leaves all results intact.
Conclusions: The CV of MTD_i determines the efficacy lost under one-size-fits-all dosing at 'the' MTD. Within plausible ranges for this CV, failure to individualize dosing can effectively halve a drug's value to society. In a competitive environment dominated by regulatory hurdles, this may reduce the value of shareholders' investment in the drug to *zero*.

Keywords

Economics of drug development, dose-finding studies, oncology, Phase I clinical trial, individualized dose-finding, precision medicine

INTRODUCTION

Dose Titration Algorithm Tuning (DTAT), a new methodology for individualized dose-finding in early-phase oncology studies, holds forth a promise of individualized dosing from the earliest stages of oncology drug development (Norris 2017). Most immediately and obviously, such individualized dosing serves the imperative of *individual ethics* in seeking to optimize the care of each person who enrolls in a Phase I study. But by increasing the efficiency of drug development overall, individualized dosing also serves wider social aims. Less effective, 'one-size-fits-all' dosing may condemn valuable drugs to failure in later registration trials. More efficacious, individualized dosing may therefore avert financial losses to shareholders in pharmaceutical innovation, while preserving innovations valuable to society at large. This brief technical note estimates the magnitude of the *social* costs incurred by one-size-fits-all dose-finding studies. The argument should be of interest to shareholders in pharmaceutical innovation, and to executives having fiduciary responsibilities to them.

THE DISTRIBUTION OF MTD_i

In (Norris 2017), DTAT was demonstrated by simulated dose titration in 25 simulated subjects drawn randomly from a population model of the pharmacokinetics and myelosuppressive dynamics of docetaxel. By extending this simulation to 1000 subjects, we obtain the empirical distribution of *individualized maximum tolerated dose* (MTD_i) shown in Figure 1.

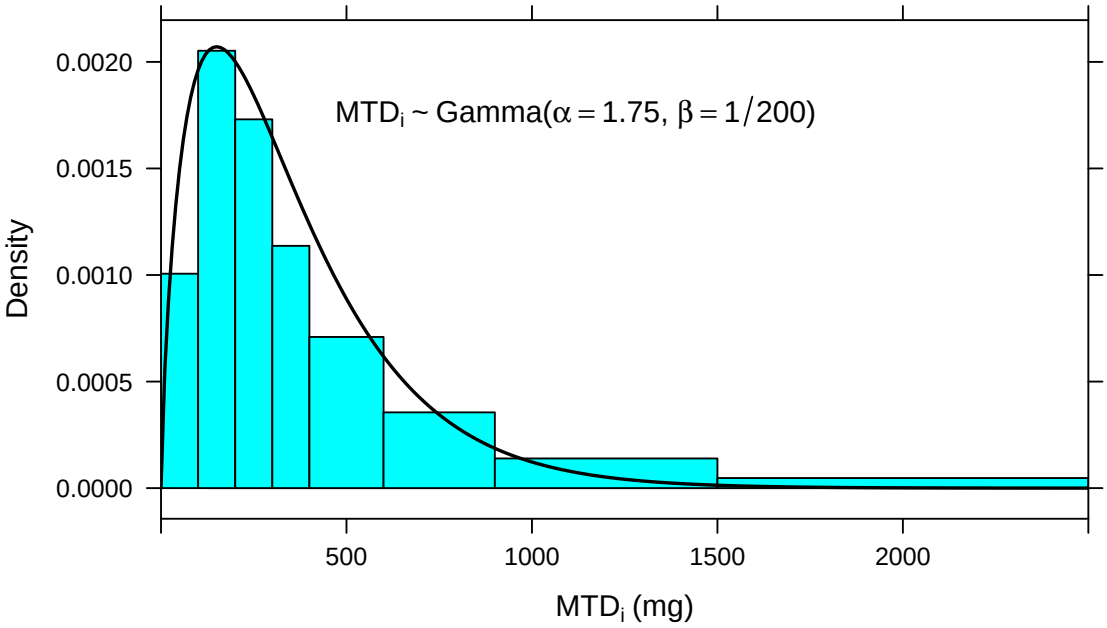

Figure 1. MTD_i is approximately Gamma distributed.

Whether the fitted Gamma density in Figure 1 represents a true distribution in any actual human population matters less for what follows than establishing the basic plausibility of a Gamma-distributed MTD_i generally.

DOSE-RESPONSE MODEL

To estimate the cost of sub-MTD_i dosing, one must model individual-level efficacy as a function of dose. A traditional approach in this context is to posit a dose-effect model of a standard 'E_{max}' type. Taking the *tumor's* point of view, we may write in fact a 'toxicology' form of the model:

$$P_r(D) = \frac{D}{D + \mathrm{LD}_{50}},$$

where $P_r(D)$ is the probability of achieving remission as a function of D, the dose received, and LD_{50} is the dose that would be 'lethal' to the tumor in 50% of patients—that is, the dose that would achieve remission with probability 0.5. By supposing further that MTD_i is the LD_{50} for the tumor in individual i, we obtain:

$$P_r = \frac{D}{D + \mathrm{MTD}_i} = \left(1 + \frac{\mathrm{MTD}_i}{D}\right)^{-1} = \left(1 + \frac{1}{\theta_i}\right)^{-1}.$$

Thus, identifying MTD_i with the LD_{50} of the tumor yields a modeled remission probability that is a function of $\theta_i = D/\mathrm{MTD}_i$, the *fraction of MTD$_i$ received*. The reasonableness of this identification will be explored in the Discussion below. As it turns out, a slightly different functional form for $P_r(\theta)$ supports obtaining an intermediate result in terms of standard functions:

$$P_r(\theta) = \frac{1}{2}\theta^{\frac{1}{2}}. \tag{1}$$

The reader suspicious of this departure from tradition should take reassurance in noting that this revised functional form is uniformly *more forgiving* of suboptimal dosing than the standard form:

$$\frac{1}{2}\theta^{\frac{1}{2}} \geq \left(1 + \frac{1}{\theta}\right)^{-1} \text{ for } \theta \geq 0.$$

THE DISTRIBUTION OF $\theta_i = \mathrm{MTD}_{\mathrm{the}}/\mathrm{MTD}_i$

If $\mathrm{MTD}_i \sim \mathrm{Gamma}(\alpha, \beta)$, then $\frac{1}{\mathrm{MTD}_i} \sim \text{Inv-Gamma}(\alpha, \beta)$ and consequently

$$\theta_i = \frac{\mathrm{MTD}_{\mathrm{the}}}{\mathrm{MTD}_i} \sim \text{Inv-Gamma}(\alpha, \beta \cdot \mathrm{MTD}_{\mathrm{the}}). \tag{2}$$

THE TWO COSTS OF ONE-SIZE-FITS-ALL DOSING

Under the prevailing practice of one-size-fits-all dosing at 'the' MTD, we take the following to occur: (1) Those individuals i for whom $\mathrm{MTD}_i > \mathrm{MTD}_{\mathrm{the}}$ will receive suboptimal dosing at a fraction $\theta_i < 1$ of their optimal dose; (2) those for whom $\mathrm{MTD}_i < \mathrm{MTD}_{\mathrm{the}}$ will experience intolerable adverse effects with a first dose, and will not receive subsequent cycles of therapy. (Those rare individuals for whom $\mathrm{MTD}_i = \mathrm{MTD}_{\mathrm{the}}$ holds *exactly* will receive their optimal $\theta_i = 1$ dose, and enjoy the full benefit of the drug.) Thus, dosing everyone at $\mathrm{MTD}_{\mathrm{the}}$ has two social costs: individuals who cannot tolerate 'the' MTD derive no benefit from the drug, while those who could have tolerated higher doses derive suboptimal benefit. This latter cost is well described in a literature stretching back 2 decades, documenting (across many types of cancer) that patients who experience milder adverse effects from chemotherapy tend to have worse outcomes (Saarto et al. 1997, Cameron et al. (2003), Di Maio et al. (2005), Yamanaka et al. (2007), Y. H. Kim et al. (2009), Lee et al. (2011), Shitara et al. (2011), McTiernan et al. (2012), Liu, Zhang, and Li (2013), Shiozawa et al. (2014), Su et al. (2015), Osorio et al. (2017)).

Figure 2 depicts the balance of these costs under 3 different choices of $\text{MTD}_{\text{the}} \in \{100, 200, 300\}$ mg. If $\text{MTD}_i \sim \text{Gamma}(\alpha = 1.75, \beta = 1/200)$, then $\theta_i = \text{MTD}_{\text{the}}/\text{MTD}_i$ will follow the inverse gamma distribution:

$$\theta_i = \frac{\text{MTD}_{\text{the}}}{\text{MTD}_i} \sim \text{Inv-Gamma}(\alpha = 1.75, \beta \in \{\frac{100}{200}, \frac{200}{200}, \frac{300}{200}\})$$

These 3 densities are plotted in green in Figure 2, superimposed on the dose-response relationship of Equation 1. Here, it is readily seen that setting $\text{MTD}_{\text{the}} = 100$ mg causes most individuals to receive doses below half of their MTD_i's ($\theta < 0.5$). Conversely, setting $\text{MTD}_{\text{the}} = 300$ mg causes few individuals to be dosed at $\theta \leq 0.5$, but excludes a large fraction of the population from treatment—as indicated by the large area under the dashed curve to the right of $\theta = 1$.

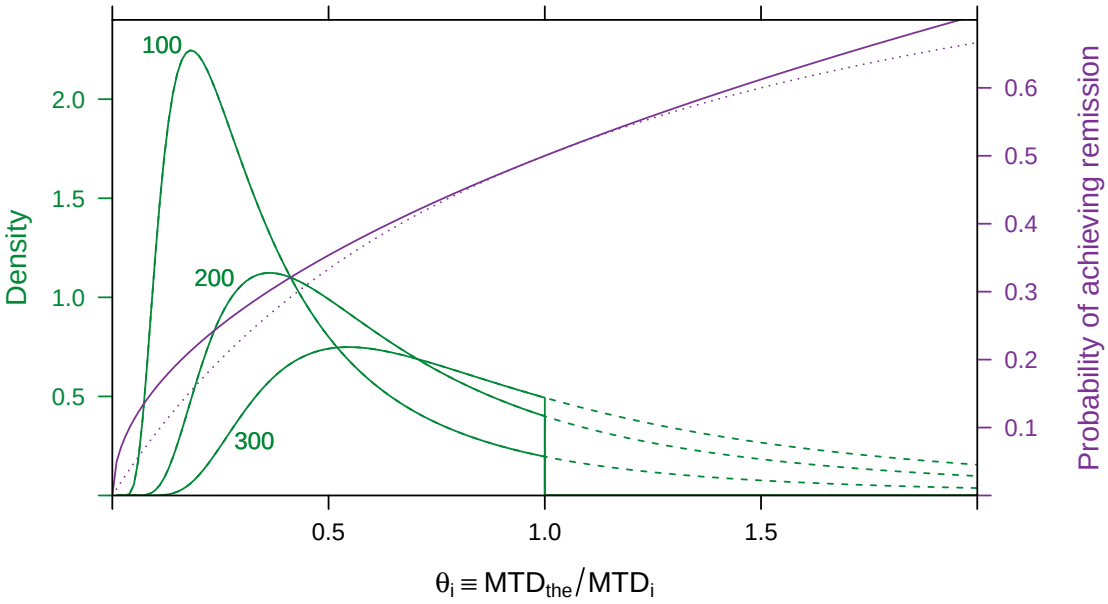

Figure 2. Social costs of one-size-fits-all dosing at 3 different choices of 'the' MTD. Against the purple dose-response function, the distribution of $\theta_i = \text{MTD}_{\text{the}}/\text{MTD}_i$ is plotted for 3 different values of MTD_{the}. When 'the' MTD is set low (100 mg), few individuals are excluded from treatment (area under dashed curves), but most are treated at a low fraction ($\theta_i < 0.5$) of their MTD_i's. Conversely, when 'the' MTD is set high (300 mg), fewer individuals are dosed so low, but many (large area under dashed curve) cannot tolerate the drug and do not receive a full course of treatment.

POPULATION-LEVEL EFFICACY OF ONE-SIZE-FITS-ALL DOSING

Given that θ_i is distributed as in Equation 2, and that the individual-level probability of remission is as given by Equation 1, then the *population rate* $\bar{P}_r$ of achieving remission may be calculated by integrating $P_r(\theta_i)$ over the treated population $0 \leq \theta_i \leq 1$. Normalizing $\tilde{\beta} = \beta \cdot \text{MTD}_{\text{the}}$, we can calculate:

$$
\begin{aligned}
\bar{P}_r &= \int_0^1 P_r(\theta) \cdot \text{Inv-Gamma}(\theta; \alpha, \tilde{\beta})\, \mathrm{d}\theta \\
&= \int_0^1 \frac{1}{2}\theta^{1/2} \cdot \frac{\tilde{\beta}^{\alpha}}{\Gamma(\alpha)} \theta^{-\alpha-1} \exp\left(-\frac{\tilde{\beta}}{\theta}\right) \mathrm{d}\theta \\
&= \frac{1}{2}\frac{\tilde{\beta}^{\frac{1}{2}}\Gamma(\alpha-\frac{1}{2})}{\Gamma(\alpha)} \int_0^1 \frac{\tilde{\beta}^{\alpha-\frac{1}{2}}}{\Gamma(\alpha-\frac{1}{2})} \theta^{-(\alpha-\frac{1}{2})-1} \exp\left(-\frac{\tilde{\beta}}{\theta}\right) \mathrm{d}\theta \\
&= \frac{1}{2}\frac{\tilde{\beta}^{\frac{1}{2}}\Gamma(\alpha-\frac{1}{2})}{\Gamma(\alpha)} \int_0^1 \text{Inv-Gamma}(\theta; \alpha-\frac{1}{2}, \tilde{\beta})\, \mathrm{d}\theta \\
&= \frac{1}{2}\frac{\tilde{\beta}^{\frac{1}{2}}\Gamma(\alpha-\frac{1}{2})}{\Gamma(\alpha)} \mathrm{Q}(\alpha-\frac{1}{2}, \tilde{\beta}),
\end{aligned}
$$

where Q denotes the regularized gamma function.

The best-case population rate of remission is obtained by choosing MTD_{the} optimally:

$$
\hat{P}_r(\alpha) = \max_{\tilde{\beta}} \left[\frac{1}{2}\frac{\tilde{\beta}^{\frac{1}{2}}\Gamma(\alpha-\frac{1}{2})}{\Gamma(\alpha)} \mathrm{Q}(\alpha-\frac{1}{2}, \tilde{\beta})\right] = \frac{1}{2}\frac{\Gamma(\alpha-\frac{1}{2})}{\Gamma(\alpha)} \max_{\tilde{\beta}} \left[\tilde{\beta}^{\frac{1}{2}} \mathrm{Q}(\alpha-\frac{1}{2}, \tilde{\beta})\right], \tag{3}
$$

in which it should be noted particularly that $\hat{P}_r$ is a function of the 'shape parameter' α, which determines the *coefficient of variation* (CV) of our gamma-distributed MTD_i via $\text{CV} = \alpha^{-1/2}$. The maximand on the right-hand side of Equation 3 is readily evaluated using the implementation of the regularized gamma function Q provided in R package **zipfR** (Evert and Baroni 2007), and the maximum obtained numerically. The dependence of $\hat{P}_r$ on $\text{CV}(\text{MTD}_i)$ is plotted in Figure 3.

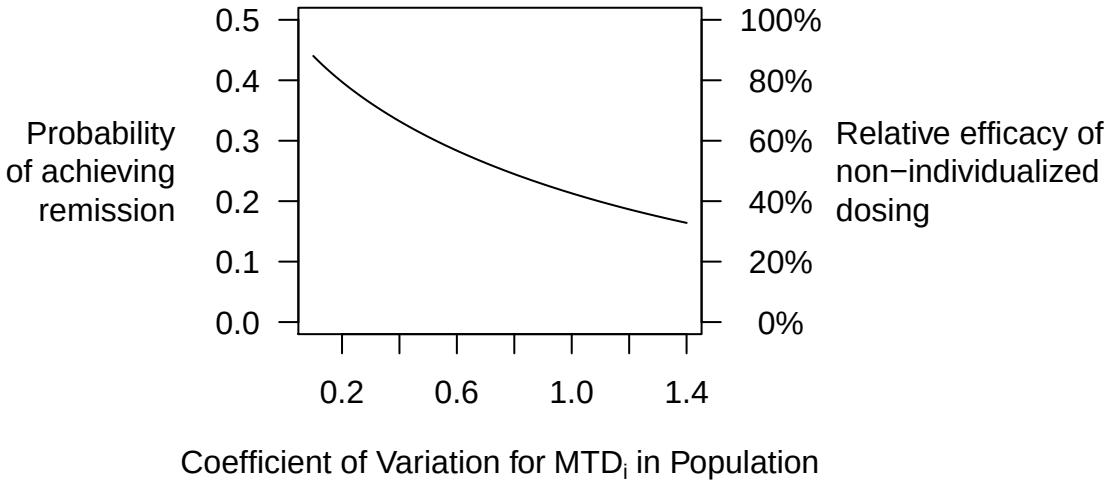

Figure 3. Estimated *minimum* cost of one-size-fits-all dosing, as a function of the coefficient of variation (CV) of MTD_i in the population. It is assumed that 'the' MTD is chosen to maximize the population-level remission rate, under the constraint of one-size-fits-all dosing. The cost of the one-size-fits-all constraint is calculated relative to a reference remission probability of 50% for optimal individualized dosing at each patient's MTD_i. The more MTD_i varies within the population, the more untenable one-size-fits-all dosing becomes.

SENSITIVITY ANALYSIS

One assumption essential to the development of my argument thus far was that individual-level outcomes are a function of $\theta_i = D/\mathrm{MTD}_i$, the fraction of MTD_i received by individual i. This assumption would hold in the limiting case where inter-individual variation in MTD_i was driven entirely by *pharmacokinetic* heterogeneity. (Consider the particularly simple example of an oral drug for which otherwise-identical individuals differed only regarding *bioavailability*.) Nevertheless, the sensitivity of my results to this assumption does seem to warrant further examination.

Consider the following perturbation of Equation 1:

$$P_r(D, \mathrm{MTD}_i) = \frac{1}{2}\left(\frac{D}{\mathrm{MTD}_i}\right)^{\frac{1}{2}}\left[1+\delta\left(\frac{\mathrm{MTD}_i}{\alpha/\beta}-1\right)\right], \tag{4}$$

where the factor in brackets induces a dependence of P_r on MTD_i that is not accounted for by $\theta_i = D/\mathrm{MTD}_i$. This factor is motivated as a first-order Taylor expansion of a general functional dependence, centered on the population mean $\mathrm{E}[\mathrm{MTD}_i] = \alpha/\beta$. Not only do we recover Equation 1 in the limit as $\delta \to 0$, but we also preserve *independently of* δ the same population-average remission rate of 1/2 under individualized 'MTD_i' dosing. (To appreciate this latter point, set $D = \mathrm{MTD}_i$ in Equation 4 then take expectations on both sides.)

To obtain $\bar{P}_r$ as previously, we rearrange Equation 4 as follows

$$P_r(D, \mathrm{MTD}_i) = \frac{1}{2}\left\{(1-\delta)\left(\frac{D}{\mathrm{MTD}_i}\right)^{\frac{1}{2}} + \delta\frac{D}{\alpha/\beta}\left(\frac{\mathrm{MTD}_i}{D}\right)^{\frac{1}{2}}\right\},$$

and then integrate as before:

$$\begin{aligned}
\bar{P}_r(\mathrm{MTD}_{\mathrm{the}}) =& \frac{1}{2}\int_0^{\mathrm{MTD}_{\mathrm{the}}} P_r(\mathrm{MTD}_{\mathrm{the}}, \mathrm{MTD}_i)\cdot \text{Inv-Gamma}(\mathrm{MTD}_i; \alpha, \beta)\,\mathrm{d}\mathrm{MTD}_i \\
=& \frac{1}{2}\int_0^1 \left[(1-\delta)\theta^{\frac{1}{2}} + \delta\frac{\mathrm{MTD}_{\mathrm{the}}}{\alpha/\beta}\theta^{-\frac{1}{2}}\right]\cdot \text{Inv-Gamma}(\theta; \alpha, \tilde{\beta})\,\mathrm{d}\theta \\
=& \frac{1}{2}\int_0^1 \left[(1-\delta)\theta^{\frac{1}{2}} + \delta\frac{\tilde{\beta}}{\alpha}\theta^{-\frac{1}{2}}\right]\cdot \text{Inv-Gamma}(\theta; \alpha, \tilde{\beta})\,\mathrm{d}\theta \\
=& \frac{1}{2}(1-\delta)\frac{\tilde{\beta}^{\frac{1}{2}}\Gamma(\alpha-\frac{1}{2})}{\Gamma(\alpha)}\,\mathrm{Q}(\alpha-\frac{1}{2}, \tilde{\beta}) + \frac{1}{2}\delta\frac{\tilde{\beta}}{\alpha}\frac{\tilde{\beta}^{-\frac{1}{2}}\Gamma(\alpha+\frac{1}{2})}{\Gamma(\alpha)}\,\mathrm{Q}(\alpha+\frac{1}{2}, \tilde{\beta}) \\
=& \frac{1}{2}\frac{\tilde{\beta}^{\frac{1}{2}}\Gamma(\alpha-\frac{1}{2})}{\Gamma(\alpha)}\,\mathrm{Q}(\alpha-\frac{1}{2}, \tilde{\beta})\left[1-\frac{\delta}{2\alpha}\right] + \frac{\tilde{\beta}^{\alpha}e^{-\tilde{\beta}}}{\Gamma(\alpha)}\left(\frac{\delta}{2\alpha}\right),
\end{aligned}$$

in which the final step requires applications of the recurrence relations for Γ and Q.

The best-case population rate of remission is obtained, as before, by choosing $\mathrm{MTD}_{\mathrm{the}}$ optimally:

$$\hat{P}_r(\alpha) = \max_{\tilde{\beta}}\left[\frac{1}{2}\frac{\tilde{\beta}^{\frac{1}{2}}\Gamma(\alpha-\frac{1}{2})}{\Gamma(\alpha)}\,\mathrm{Q}(\alpha-\frac{1}{2}, \tilde{\beta})\left[1-\frac{\delta}{2\alpha}\right] + \frac{\tilde{\beta}^{\alpha}e^{-\tilde{\beta}}}{\Gamma(\alpha)}\left(\frac{\delta}{2\alpha}\right)\right]. \tag{5}$$

Equation 5 generalizes Equation 3, and reduces to it on setting $\delta = 0$.

Without attempting to develop an interpretation of δ, we plot in Figure 4 the relative efficacy of one-size-fits-all dosing at selected values of CV, for $\delta \in [-1, 1]$. Evidently, the perturbation is of second order. Moreover, absolute deviations from the picture of Figure 3 prove small over the range of what seems from Equation 4 to be the natural scale for δ.

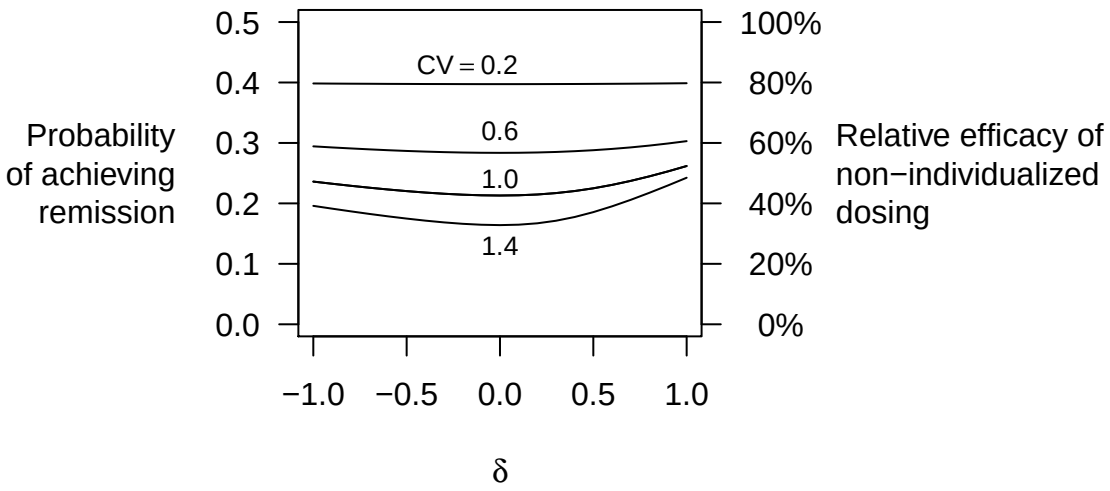

Figure 4. Sensitivity of efficiency loss estimate to perturbation as in Equation 4.

Sensitivity under an interpretable alternative

As an alternative to the foregoing perturbation analysis, we might instead posit a single, readily interpretable alternative to Equation 1. A possibility that immediately presents itself as a 'diametrically opposed' alternative is to suppose $P_r(D, \mathrm{MTD}_i)$ *independent of* MTD_i, as in:

$$P_r^*(D, \mathrm{MTD}_i) \equiv P_r^*(D) = \gamma \left(\frac{D}{\alpha/\beta} \right)^{\frac{1}{2}} = \gamma \left(\frac{D}{\mathrm{E}[\mathrm{MTD}_i]} \right)^{\frac{1}{2}}. \tag{6}$$

Under one-size-fits-all dosing, Equation 6 yields a population-level efficacy of

$$\bar{P}_r^*(D) = \int_D^\infty P_r^*(D) \cdot \mathrm{Gamma}(x; \alpha, \beta)\, \mathrm{d}x = \gamma \left(\frac{D}{\alpha/\beta} \right)^{\frac{1}{2}} \mathrm{Q}(\alpha, \beta \cdot D). \tag{7}$$

In contrast to the cases considered above, here the integration extends into a region where $P_r > 1$; but, by taking $\gamma \to 0$, we can push this region as far as desired into the tail of our Gamma-distributed MTD_i. (As will be seen presently, γ drops out of our analysis of *relative* efficacy.) Differentiating with respect to D, and employing the same normalization $\tilde{\beta} = \beta \cdot D$ introduced previously, we find that Equation 7 is maximized at $\tilde{\beta} = \hat{\beta}$ defined by

$$\mathrm{Q}(\alpha, \hat{\beta}) = 2 \frac{\hat{\beta}^\alpha e^{-\hat{\beta}}}{\Gamma(\alpha)}. \tag{8}$$

Substituting (8) into (7), we find that this maximum value is:

$$\hat{P}_r^* = \gamma \left(\frac{\hat{\beta}}{\alpha} \right)^{\frac{1}{2}} 2 \frac{\hat{\beta}^\alpha e^{-\hat{\beta}}}{\Gamma(\alpha)} = 2\gamma \frac{\hat{\beta}^{\alpha + \frac{1}{2}} e^{-\hat{\beta}}}{\sqrt{\alpha}\Gamma(\alpha)}. \tag{9}$$

Under the MTD_i-independent remission probability (6), optimal individualized dosing does not reduce trivially to a constant 1/2. Rather, we must calculate as follows:

$$\bar{P_r}^{*,i} = \int_0^\infty P_r^*(D) \cdot \mathrm{Gamma}(D;\alpha,\beta)\,\mathrm{d}D = \gamma \frac{\Gamma(\alpha+\frac{1}{2})}{\sqrt{\alpha}\Gamma(\alpha)}. \tag{10}$$

Dividing (9) by (10), we obtain the following expression for relative efficacy of one-size-fits-all dosing under Equation 6:

$$\frac{\hat{P_r^*}}{\bar{P_r}^{*,i}} = 2\frac{\hat{\beta}^{\alpha+\frac{1}{2}}e^{-\hat{\beta}}}{\Gamma(\alpha+\frac{1}{2})}. \tag{11}$$

Solving Equation 8 numerically for $\hat{\beta} = \hat{\beta}(\alpha)$, and substituting this into Equation 11, we obtain the following counterpart to Figure 3:

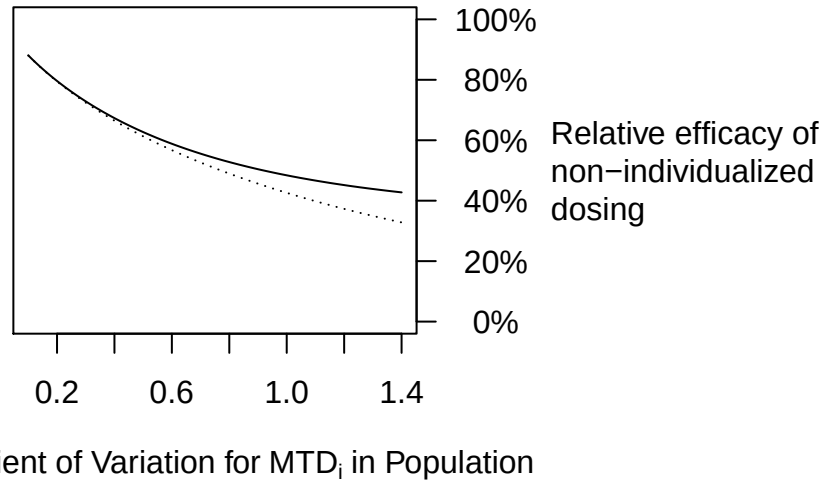

Figure 5. Sensitivity of Figure 3 to alternative specification of an MTD_i-independent dose-effect relation as in Equation 6. The solid curve shows efficacy loss under this alternative specification, while the dotted curve shows, for comparison, the original curve derived from Equation 1. Evidently, this alternative specification does little to redeem one-size-fits-all dosing.

DISCUSSION

At two points in this argument, I have adopted modeling assumptions that are relatively *forgiving* of one-size-fits-all dosing, and therefore would tend to underestimate its costs. Firstly, my highly concave square-root $\mathrm{E_{max}}$ model (Equation 1) regards under-dosing more favorably than does a typical $\mathrm{E_{max}}$ model, such as appears in the dotted purple curve in Figure 2. Secondly, the optimization itself in Equation 3 surely overestimates the population-level outcomes achieved by one-size-fits-all dosing as implemented in current Phase I designs. Indeed, these designs tend to target DLT *rates* without explicit reference or regard to outcomes.

Dose reduction protocols, as seen both in trials and in clinical practice, do somewhat relax the extreme form of one-size-fits-all dosing constraint that I have modeled in this paper. Clearly, such protocols exist precisely to recover some part of the lost value I calculate here. But given that these protocols are readily interpreted as a (very) poor man's DTAT, their existence only underscores the urgent need for rational dose individualization in oncology. This conclusion fully withstands the rather vigorous sensitivity analysis performed above.

CONCLUSIONS

Taking *population-level efficacy* as a proxy, I have estimated the social cost of one-size-fits-all dosing organized around a concept of 'the' maximum tolerated dose (MTD) in oncology. The magnitude of this cost is seen to depend primarily on the coefficient of variation (CV) of *individually optimal* MTD_i doses in the population. Within plausible ranges for this CV, the failure to individualize dosing can effectively halve a drug's value to society. Notably, in a competitive environment dominated by regulatory hurdles, this may reduce the value of shareholders' investment in a drug to *zero*.

DATA AVAILABILITY

Open Science Framework: Data for Figure 1 may be found in R package **DTAT** (v0.1-1), available together with code for reproducing all of this paper's Figures and analyses, at doi: 10.17605/osf.io/vtxwq.

Competing interests

The author operates a scientific and statistical consultancy focused on precision-medicine methodologies such as those advanced in this article.

Grant information

The author declared that no grants were involved in supporting this work.

REFERENCES

Cameron, D. A., C. Massie, G. Kerr, and R. C. F. Leonard. 2003. "Moderate Neutropenia with Adjuvant CMF Confers Improved Survival in Early Breast Cancer." *Br. J. Cancer* 89 (10): 1837–42. doi:10.1038/sj.bjc.6601366.

Di Maio, Massimo, Cesare Gridelli, Ciro Gallo, Frances Shepherd, Franco Vito Piantedosi, Silvio Cigolari, Luigi Manzione, et al. 2005. "Chemotherapy-Induced Neutropenia and Treatment Efficacy in Advanced Non-Small-Cell Lung Cancer: A Pooled Analysis of Three Randomised Trials." *Lancet Oncol.* 6 (9): 669–77. doi:10.1016/S1470-2045(05)70255-2.

Evert, Stefan, and Marco Baroni. 2007. "*zipfR*: Word Frequency Distributions in R." In *Proceedings of the 45th Annual Meeting of the Association for Computational Linguistics, Posters and Demonstrations Sessions*, 29–32. Prague, Czech Republic.

Kim, Yun Hwan, Hyun Hoon Chung, Jae Weon Kim, Noh-Hyun Park, Yong-Sang Song, and Soon-Beom Kang. 2009. "Prognostic Significance of Neutropenia During Adjuvant Concurrent Chemoradiotherapy in Early Cervical Cancer." *J Gynecol Oncol* 20 (3): 146–50. doi:10.3802/jgo.2009.20.3.146.

Lee, C. K., H. Gurney, C. Brown, R. Sorio, N. Donadello, G. Tulunay, W. Meier, et al. 2011. "Carboplatin-Paclitaxel-Induced Leukopenia and Neuropathy Predict Progression-Free Survival in Recurrent Ovarian Cancer." *Br. J. Cancer* 105 (3): 360–65. doi:10.1038/bjc.2011.256.

Liu, Wei, Cui-Cui Zhang, and Kai Li. 2013. "Prognostic Value of Chemotherapy-Induced Leukopenia in Small-Cell Lung Cancer." *Cancer Biol Med* 10 (2): 92–98. doi:10.7497/j.issn.2095-3941.2013.02.005.

McTiernan, Anne, Rachel C. Jinks, Matthew R. Sydes, Barbara Uscinska, Jane M. Hook, Martine van Glabbeke, Vivien Bramwell, et al. 2012. "Presence of Chemotherapy-Induced Toxicity Predicts Improved Survival in Patients with Localised Extremity Osteosarcoma Treated with Doxorubicin and Cisplatin: A Report from the European Osteosarcoma Intergroup." *Eur. J. Cancer* 48 (5): 703–12. doi:10.1016/j.ejca.2011.09.012.

Norris, David C. 2017. "Dose Titration Algorithm Tuning (DTAT) Should Supersede 'the' Maximum Tolerated Dose (MTD) in Oncology Dose-Finding Trials." *F1000Research* 6 (March): 112. doi:10.12688/f1000research.10624.2.

Osorio, J. C., A. Ni, J. E. Chaft, R. Pollina, M. K. Kasler, D. Stephens, C. Rodriguez, et al. 2017. "Antibody-Mediated Thyroid Dysfunction During T-Cell Checkpoint Blockade in Patients with Non-Small-Cell Lung Cancer." *Ann Oncol*. Accessed March 6. doi:10.1093/annonc/mdw640.

Saarto, T., C. Blomqvist, P. Rissanen, A. Auvinen, and I. Elomaa. 1997. "Haematological Toxicity: A Marker of Adjuvant Chemotherapy Efficacy in Stage II and III Breast Cancer." *Br. J. Cancer* 75 (2): 301–5.

Shiozawa, Yusuke, Junko Takita, Motohiro Kato, Manabu Sotomatsu, Katsuyoshi Koh, Kohmei Ida, and Yasuhide Hayashi. 2014. "Prognostic Significance of Leukopenia in Childhood Acute Lymphoblastic Leukemia." *Oncol Lett* 7 (4): 1169–74. doi:10.3892/ol.2014.1822.

Shitara, Kohei, Keitaro Matsuo, Isao Oze, Ayako Mizota, Chihiro Kondo, Motoo Nomura, Tomoya Yokota, Daisuke Takahari, Takashi Ura, and Kei Muro. 2011. "Meta-Analysis of Neutropenia or Leukopenia as a Prognostic Factor in Patients with Malignant Disease Undergoing Chemotherapy." *Cancer Chemother Pharmacol* 68 (2): 301–7. doi:10.1007/s00280-010-1487-6.

Su, Zhen, Yan-Ping Mao, Pu-Yun OuYang, Jie Tang, Xiao-Wen Lan, and Fang-Yun Xie. 2015. "Leucopenia and Treatment Efficacy in Advanced Nasopharyngeal Carcinoma." *BMC Cancer* 15 (May): 429. doi:10.1186/s12885-015-1442-3.

Yamanaka, T., S. Matsumoto, S. Teramukai, R. Ishiwata, Y. Nagai, and M. Fukushima. 2007. "Predictive Value of Chemotherapy-Induced Neutropenia for the Efficacy of Oral Fluoropyrimidine S-1 in Advanced Gastric Carcinoma." *Br J Cancer* 97 (1): 37–42. doi:10.1038/sj.bjc.6603831.