Received: 30 September 2021 | Revised: 20 February 2022 | Accepted: 6 March 2022

DOI: 10.1111/bcp.15307

ORIGINAL ARTICLE

Dose banding – Weighing up benefits, risks and therapeutic failure

Stephen Duffull

School of Pharmacy, University of Otago, Dunedin, New Zealand

Correspondence
Stephen Duffull, School of Pharmacy, University of Otago. Dunedin, New Zealand.
Email: stephen.duffull@otago.ac.nz

[Correction added on 19 May 2022, after first online publication: CAUL funding statement has been added.]

Aims: Dose banding is a method of dose individualisation in which all patients with similar characteristics are allocated to the same dose. Dose banding results in some patients receiving less intensive treatment which risks a reduction in therapeutic benefit (iatrogenic therapeutic failure) because of variability not predicted by dose banding. This study aims to explore the effects of dose banding on therapeutic success and failure.

Methods: This was a simulation study. Virtual patients were simulated under a simple pharmacokinetic model where the response of interest is the steady-state average concentration. Clearance was correlated with a covariate used for dose banding. Dose individualisation was based on: one-dose-fits-all, covariate-based dosing, empirical dose banding, dose banding optimised for net therapeutic benefit and optimised for both benefit and minimising iatrogenic therapeutic failure.

Results: The lowest and highest probability of target attainment (PTA) were 44% for one-dose-fits-all and 72% for covariate-based dosing. Neither dosing approach would result in iatrogenic therapeutic failure as lower dose intensities do not occur. Empirical dose banding performed better than one-dose-fits-all with 59% PTA but not as good as either optimised method (64–69% PTA) while carrying a risk of iatrogenic therapeutic failure in 25% of patients. Optimising for benefit (only) improved PTA but carried a risk of iatrogenic therapeutic failure of up to 10%. Optimising for benefit and minimising iatrogenic therapeutic failure provided the best balance.

Conclusion: Future application of dose banding needs to consider both the probability of benefit as well the risk of causing iatrogenic therapeutic failure.

KEYWORDS
bioethics, dose banding, dose individualisation, pharmacometrics

1 | INTRODUCTION

The purpose of dose individualisation is to optimise clinical outcomes for patients who fall in the tails of the population, i.e., those who have a much greater or more attenuated drug response compared to the standard person[1] (see Figure 1 [upper panels]). Standard guidelines and clinical pathways are designed to optimise treatment for the standard patient but cannot, by definition, help those who are non-standard (for instance the patient who has threefold higher or lower dose requirement than predicted for a standard patient). The example shown in the upper right panel of Figure 1 illustrates a single dose level optimised for the standard patient which then results in highly

variable responses between (non-standard) patients. Dose individualisation, whether at the level of the starting dose or via dose titration based on a measured response will help individualise the response of the patient to their target clinical response.[2] The example shown in Figure 1 (lower panels) illustrates perfect dose individualisation (each patient's dose is based on their true parameter values). This level of dose individualisation is impossible to achieve clinically (at best we can individualise based on the covariate only) but provides a useful positive control.

Dose banding is a commonly used form of individualisation of the starting dose and is characterised by allocating patients to a pre-specified dose group based on a particular observable characteristic.[3–5] The characteristic may be related to drug clearance, such as renal function, or it may relate to a diagnostic label (e.g., disease staging) or perhaps some other (arbitrary) feature. In the former case the characteristic will almost certainly be a covariate identified from a population pharmacokinetic model and hence is assumed to be on the causal path from dose to some measure of exposure. In the case of a diagnostic label, this may define the type or intensity of a treatment based on some prognostic indicator. We may also have a mixture of the two, where a dose band is formed based on a covariate (e.g., renal function) but the bands themselves are based on diagnostic labels (e.g., chronic kidney disease [CKD] diagnostic bands 2, 3a, 3b, etc.). This composite approach is appealing in its ease of uptake by clinical staff but is not optimised for choice of dose since the bands were developed for diagnostic-prognostic purposes rather than optimising treatment intensity. Some common examples of dose banding include: dosing guidelines for paracetamol based on age, dosing guidelines for metformin based on renal function, and dosing guidelines for dabigatran based on age and renal function. The various guidelines can be reviewed in their respective drug labels.

What is already known about this subject

- Dose banding is a widely used and accepted method of dose individualisation. Its simplicity and wide applicability make it an attractive dosing tool. Optimising dose banding has been shown to improve therapeutic benefits.

What this study adds

- This study addresses the implications of being allocated to a lower dose-rate (i.e., a reduction dose level or increase in interval) due to dose banding. While dose rate is maintained in some patients, others may have a reduction leading to the potential risk of therapeutic failure. Methods are introduced to reduce this risk while maximising therapeutic benefits.

The simplest form of dose banding involves the allocation of all patients to receive the same dose irrespective of their characteristics (e.g., the dose of paracetamol in adults). This method is essentially a *one-dose-fits-all* approach which is appealing in its simplicity but lacks granularity that may be needed for other medicines in order to achieve optimal clinical outcomes for all. The corollary to one-dose-fits-all is where each patient is essentially their own dose group (e.g., a continuous dose adjustment, such as mg/kg of body weight—as used, for example, with infliximab). This is conceptually *infinite dose banding* limited only by the resolution available in the dosage formulation or measurement of the observable characteristic. Since, however, we do

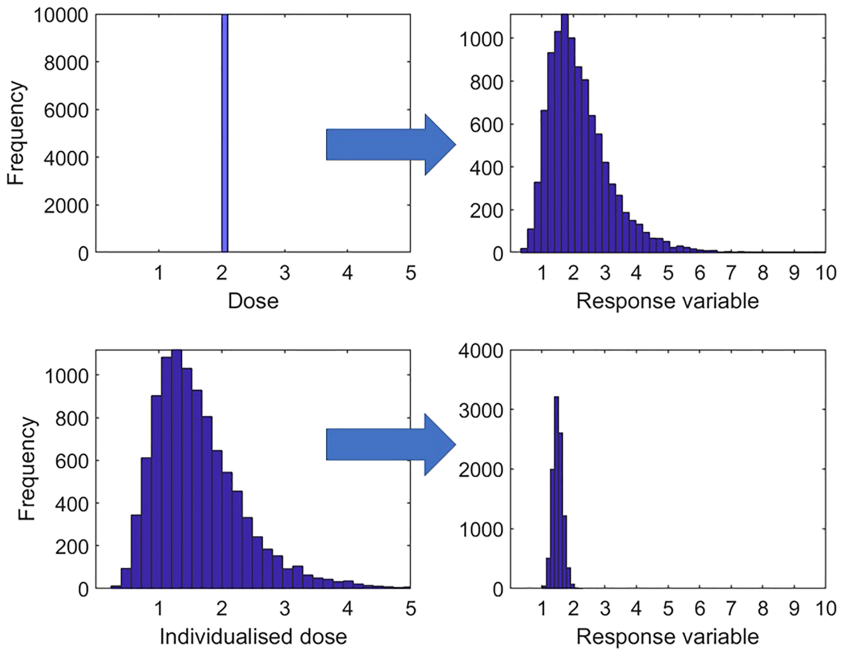

FIGURE 1 The left panels represent the frequency of the virtual population *vs* the dose that they received. The right panels show the frequency distribution of the response variable which is a function of dose, variability in clearance and variability in the residual error (as per Equations 1–5). A single dose level (upper right panel) leads to a wide range of values of the response variable (upper left panel). Our goal is to reduce the range (tails) of the response variable so patients experience a more consistent response. If we apply perfect individualisation of each patient's dose (lower left panel) then the range of doses used across the population increases dramatically while reducing the tails of the distribution of patient response (lower right panel) which now lies within the desirable range

not have access to infinite dosage formulations, this approach is termed *fully individualised*.

The most common approach to dose banding lies between these two extremes in which two (or more) bands are defined and patients are allocated to a dose band based on their characteristic(s). A common clinical instance of dose banding occurs based on creatinine clearance (CLCR) but can also be based on cost,[3] available formulation size[6] or doses used in a fixed dose combination.[1] For instance, a dose band often occurs at pre-specified creatinine clearance values,[7] say for example 30 mL/min. In this example (as described by Doogue and Polasek[7]), a patient who has a value of creatinine clearance of 29 mL/min would receive a lower dose ($D^{-}_{CLCr<30}$) compared to a patient who had a creatinine clearance of 30 mL/min ($D^{+}_{CLCR\geq 30}$). The notation D^{+} represents preservation of dose intensity for patients at or above the cut-score (in this case 30 mL/min) and D^{-} a reduction in dose intensity for those below the cut-score. Here dose intensity has the same meaning as dose rate and depends on the dose level as well as the interval, and increasing dose intensity reflects an increase in dose rate. A reduction in dose intensity may risk therapeutic failure. Of course, measurement of serum creatinine is subject to both circadian variability[8] and significant intra- and interlaboratory assay error[9] as well as known issues with equations used to approximate CLCR,[10] which risks misclassification of patients into dose bands based on variation in measurement of creatinine clearance rather than variation in glomerular filtration rate.

The purpose of dose banding is to provide simple guidelines for initiating treatment that maximises the probability of success (doing good = *beneficence*) while minimising harm due to side effects (causing harm = *maleficence*). An excellent description of a methodology to achieve optimal dose banding for therapeutic benefit is provided in Jönsson and Karlsson.[5] In the majority of cases, however, dose individualisation is not focused on improving beneficence but rather reducing maleficence, which is the case for metformin, where dose banding is attempting to reduce the chance of rare severe acidotic events.[11] Recent work further highlights issues with dose banding in that our desire to do good (or more likely to avoid doing bad) causes us not to do good. Here, this is termed *non-beneficence* which is used to refer to opportunity loss caused by withholding the more beneficent treatment approach. There are examples of non-beneficence at play. The work by Wright et al.[12] identified that a cut-score of creatinine clearance of 40 mL/min, a covariate for allopurinol dosing, resulted in a > 60% increased risk of therapeutic failure in the $D^{-}_{CLCR<40}$ group compared with those that continued dose intensity. Similarly, the diagnostic banding from the work of Venkataramani et al.[13] illustrated that a CD4+ value of ≥ 200/mL was associated with 35% higher mortality than a CD4+ value of 199 or lower on the basis that this cut-score was used for choice of treatment. Those with a CD4+ count of ≥ 200/mL were not offered treatment.

From a bioethical standpoint and for the purposes of this study, it is considered that beneficence, the act of doing good, has a corollary "non-beneficence", the act of not doing good. In the description here, non-beneficence is distinguished from maleficence by virtue of process. Non-beneficence denotes the failure to act and therefore the failure to create benefit (an opportunity loss) and maleficence the process of action in which a harmful event occurs. Note some diversity exists about whether harm by omission constitutes maleficence or lack of beneficence.[14] Since the concept is central to this work, the distinction between non-beneficence and maleficence is made because (1) it is evident in dose individualisation where harm is considered to occur at high concentrations but is often ignored at low concentrations, (2) non-beneficence includes potential and actual harm from opportunity loss, and (3) both of these concepts can be distinguished from harm due to excessive dosing (maleficence). This distinction aligns closely with the work on the effects of cut-scores for diagnostic reasoning.[15]

The benefit of dose banding is the simplicity for adjusting doses based on a patient's characteristics. There are two main issues with dose banding: (1) dose banding results in loss of granularity in dose intensity by applying a discrete set of doses which may reduce the probability of success rather than selecting from a continuous dose range and (2) patients may be disadvantaged if they have a set of characteristics that are just above or below the dose banding cut-off such that they receive an increased or reduced dose intensity. Of note, dose banding is not applied to a measured drug response which would fall in the realm of a nomogram or adaptive Bayesian forecasting strategy.

This study aims to explore the attributes of dose banding in relation to beneficence, maleficence and non-beneficence. The specific objectives were to explore the operating characteristics of: (i) one-dose-fits-all, (ii) covariate-based dosing (fully individualised starting dosing equivalent to group dosing[16]), (iii) fixed dose levels with dose banding, and (iv) optimised dose levels and optimised banding cut-scores.

2 | METHODS

All aims were addressed by simulation and optimisation using MATLAB version 9.9.0.1467703 (R2020b). In this study the simplest pharmacokinetic model was considered in order to explore the potential influence of dose banding. All simulations included 10 000 virtual patients.

2.1 | Application of dose banding

Dose banding defines a range of a patient characteristic within which all patients are given the same dose. For the one-dose-fits-all category, there is no cut-score and hence there is only one dose (D) for all simulated subjects. For two bands, the cut-score (termed k) is the value of a characteristic of interest (z) at which the dose of the drug would be changed. If the patient's value of their characteristic z is lower than k then the dose intensity is given by $D^{-}_{z<k}$ (the lower dose); if it is equal to or higher than k then the current dose intensity $D^{+}_{z\geq k}$ is continued. For example, if the value of creatinine clearance is <30 mL/min then the dose of dabigatran would be reduced. The choice of

TABLE 1 Parameter values used in the simulation

Parameter[a]	Value
$\mu_{CL}\ (L/d)$	1
$\omega_{CL}\ (CV\%)$	30%
μ_z	1
$\omega_z\ (CV\%)$	30%
$\sigma_1\ (CV\%)$	10%
$\sigma_2\ (mg/L)$	0.01

[a]Units are arbitrary.

dose level is defined by comparing z (the covariate) to k (the cut-score). This can be easily extended to $n-$ dose bands where the number of dose bands equals the number of cut-score values + 1.

2.2 | Simulation models

The simplest pharmacokinetic model was used to explore dose banding. The simulation was based on a one-parameter steady-state model defined only by the parameter clearance (CL) and variable dose rate. CL was positively (linearly) correlated with an influential covariate, denoted z. Equation 1 depicts the covariate distribution model for z, Equation 2 the between-subject variability on clearance (that is unexplained by the covariate z), Equation 3 the fixed effects model for clearance and Equations 4 and 5 the model for the response variable and residual error, respectively.

$$\ln(z_i) \sim N(\ln\ (\mu_z), \omega_z) \tag{1}$$

$$\eta_i \sim N(0,\ \omega_{CL}) \tag{2}$$

$$CL_i = \mu_{CL} \cdot \frac{z_i}{mean(z)} \cdot \exp(\eta_i) \tag{3}$$

$$c_{ss,ave(i)} = \frac{Dose_i}{CL_i} \cdot \exp(\varepsilon_{1_i}) + \varepsilon_{2_i} \tag{4}$$

$$\begin{bmatrix} \varepsilon_1 \\ \varepsilon_2 \end{bmatrix} \sim N\left(0, \begin{bmatrix} \sigma_1 & 0 \\ 0 & \sigma_2 \end{bmatrix}\right) \tag{5}$$

where μ_z and μ_{CL} are the population mean values of the covariate z and clearance CL and ω_z and ω_{CL} the standard deviations of the between-subject variability in z and CL, respectively, $mean(z)$ is μ_z, η_i is the difference of the i^{th} individual's value of CL from the population mean value and CL_i is the i^{th} individual's value of clearance, $c_{ss,ave(i)}$ represents the i^{th} simulated subject's average steady-state concentration based on the dose rate $(Dose_i)$, and $\varepsilon_{1\text{ or }2}$ are residual differences (in this case proportional and additive) with standard deviations $\sigma_{1\text{ or }2}$. The units are arbitrary. In this notation, ~ is used to signify distributed as or when simulating that the random variable was drawn from the distribution illustrated on the right-hand side. No measurement error was included for z.

While this model may seem overly simplistic, for chronic dosing (to steady state) we may be only interested in exposure (area under the concentration–time curve), which is proportional to $c_{ss,ave}$.

The parameter values are given in Table 1.

2.3 | Determination of a utility for beneficence, maleficence and non-beneficence

The benefit or otherwise of a treatment was based on a therapeutic range defined by a lower and upper boundary value (equivalent to a square loss function). The success of a treatment was defined according to the three cases in Equation 6. The upper case corresponds to the situation in which values of $c_{ss,ave}$ are within the range and hence is a success. For the purposes of the benefit and risk assessment, it was assumed that in all cases when the $C_{ss,ave}$ concentration was between $1-2$ mg/L that the benefits always outweighed the risks and hence benefits and risks were not separately calculated but rather considered as a joint concept. The middle case, where the virtual patient has a simulated value of $c_{ss,ave}$ below the lower bound, is considered a potential therapeutic failure, and the lower case, where the concentration is above the upper bound, risks toxicity (i.e., where risks outweigh the benefits). The success (s) from any given dosing regimen for the i^{th} patient was therefore:

$$\begin{cases} L_{c_{ss,ave}} \le c_{ss,ave,(i)} \le U_{c_{ss,ave}}, & s_i = 1\ \text{beneficence (success)} \\ c_{ss,ave,(i)} < L_{c_{ss,ave}}, & s_i = 0\ \text{therapeutic failure (see later descrption)} \\ U_{c_{ss,ave}} > c_{ss,ave(i)}, & s_i = 0\ \text{maleficence (harm from excess dose)} \end{cases} \tag{6}$$

In this simulation, a net benefit which equates to a success of treatment and assigned the value of 1, was determined when the average steady-state concentrations ($c_{ss,ave}$) fell between 1 (L_{Css}) and 2 (U_{Css}) mg/L. If $c_{ss,ave} > 2$ this represents a net increased risk of toxicity (potential maleficence), assigned a value of 0, and if $c_{ss,ave} < 1$ a net increased risk of therapeutic failure, assigned a value of 0. For this criterion, the cause of the therapeutic failure is not relevant for the purposes of determining success. The empirical probability of success (probability of target attainment; PTA) for a simulated cohort of patients was given by the number of successes as a proportion of the total number of virtual patients.

$$\text{PTA}_1 = \frac{1}{m}\sum_{i=1}^{m} s_i \tag{7}$$

and m is the number of patients in the simulation. The probability of target attainment when just the net benefit situation is considered as per Equation 7 is termed PTA_1.

Therapeutic failure, where $c_{ss,ave} < 1$ mg/L, was further divided into treatment failure (random failure) where $Dose = D_k^+$ (i.e., not adjusted to a lower band) and non-beneficence (iatrogenic, failure to intensify treatment) where the dose has been reduced according to $Dose = D_k^-$ for a given cut-score (k). Note, iatrogenic failure, in this context, does not imply that the prescriber has made an error but

rather that the guidelines were not optimal. The probability of success from Equation 7 was modified (for some simulation scenarios) with an additional penalty for non-beneficence. The outcome of a failed treatment was divided into non-beneficence and random failure and is shown in Equation 8 with the two corresponding cases. The upper case describes the situation in which the concentration was below the lower bound *and* (using the symbol ∩) the dose intensity had been reduced as $z <$ cut-score. This represents a non-beneficent case setting (i.e., failure to achieve the target because the dose had been reduced). The lower case is all other settings in which therapeutic failure due to non-beneficence did not occur. Note that it is assumed here that the dose adjustment is always correctly changed in relation to the patient's value of their covariate z. Hence, iatrogenic failure (non-beneficence) is caused because of following the dose-banding guidelines and not because of prescriber error.

$$\begin{cases} c_{ss,ave(i)} < L_{c_{ss,ave}} \cap Dose = D_k^-, & \upsilon = 1 \\ \text{otherwise}, & \upsilon = 0 \end{cases} \quad (8)$$

The probability of target attainment is now a function of both beneficent success and not non-beneficent failure. This is illustrated in Equation 9 where the first term is the same as in Equation 7 and the second denotes the not non-beneficent situation.

$$\text{PTA}_2 = \left(\frac{1}{m}\sum_{i=1}^{m} s_i\right)\left(1 - \frac{1}{m}\sum_{i=1}^{m} \upsilon_i\right) \quad (9)$$

The probability of target attainment when the compositive of both net benefit and non-beneficent were considered as per Equation 9 is termed PTA_2. The second term represents 1− the probability of non-beneficence.

2.4 | Optimisation of dose banding dose rates and cut-scores

Optimisation of both dose rate(s) and cut-score(s) of z were performed simultaneously using simulated annealing. The search was based on optimising the PTA (either PTA_1 [Equation 7] or PTA_2 [Equation 9] depending on the simulation scenario) for a cohort of 10 000 virtual patients, where larger values of PTA are better. The virtual patients were simulated according to Equations 1–5 and the $PTA_{1\text{ or }2}$ was calculated over the 10 000 patients based on the current values of the dose rate (as appropriate for the patient's covariate value). The value of the cut-score(s) of z and dose rate for each dose band were selected by the search algorithm that maximised these PTA values (from either Equations 7 or 9) across this virtual cohort.

2.5 | Simulation-optimisation scenarios

Twelve simulation +/− optimisation scenarios were considered (details are shown in Table 2). These included five fixed doses with dose bands (F1 … F4, FI) based on a priori knowledge of the system, four optimised doses and cut-scores (O1 … O4) under a standard net benefit setting (Equation 7) and three optimised doses and cut-scores (ONB2 … ONB4) for the composite of both net benefit and non-beneficence (Equation 9). ONB1 was not considered as when there is only one dose band then there can be no non-beneficence since everyone gets the same dose intensity. In addition, optimised dose banding was not considered for fully individualised dosing.

The setting with a single dose level is equivalent to the one-dose-fits-all scenario and fully individualised is equivalent to a full covariate-based dosing regimen (e.g., every virtual patient gets their own unique dose given their own unique covariate value).

The fixed dose levels and cut-scores were based on the percentiles of the distribution of z. For a single dose level (one-dose-fits-all), the dose was set to 2 mg/d (there are no cut-scores for a single dose level). For the two-dose level setting, the cut-score was at the 50th percentile for z and the doses were 2 mg/d and 1 mg/d if the value of z for a virtual patient was higher or lower than the cut-score, respectively. For the three-dose level setting, the cut-scores were at the 67th and 33rd percentiles of z and the doses were 2 mg/d, 1.34 mg/d and 0.66 mg/d (respectively).

3 | RESULTS

In this simulation study the probability of target attainment (PTA), i.e. the probability that the value of $c_{ss,ave}$ was between 1 and 2 mg/L, was evaluated under different dose banding settings as well as the probability of a non-beneficent outcome. Non-beneficence was

TABLE 2 Simulation scenarios

Number of dose levels[a]	Fixed dose bands	Optimised dose bands net benefit risk only	Optimised dose bands for net benefit and non-beneficence
1	F1	O1	–
2	F2	O2	ONB2
3	F3	O3	ONB3
4	F4	O4	ONB4
Fully individualised	FI	–	–

[a]The number of cut-scores = number of dose levels minus 1.

defined as a potential therapeutic failure (evaluated as a $c_{ss,ave} < 1$ mg/L) caused by assigning a patient to a lower dose level (D^-) based on dose-banding rules. The results are presented in Table 3. The simplest setting has only one dose level, one-dose-fits-all, and hence has no dose banding and therefore avoids the possibility of non-beneficent outcomes. In this setting we see the probability of target attainment is 44% for a fixed dose (2 mg/d) setting and 58% for an optimised dose (1.4 mg/d). Increasing the between-subject variability of *CL* to 50% decreased the PTA to 37% and 45% for the fixed dose and optimised dose respectively (results are not shown in Table 3). The best-case scenario is given by a covariate-based dosing, where every patient gets a dose individualised to their personal covariate values (termed fully individualised dosing). Here the PTA was 0.72. Improvement past this value would require dose titration to the target (e.g., using target concentration intervention or similar). The distribution of $c_{ss,ave}$ responses for the 10 000

TABLE 3 Results of different dose banding scenarios

No. dose levels	Method	PTA[a]	Cut-score (percentiles of *z*)	Dose levels (*mg*/*d*)	PrNB
	Fixed	0.44	n/a	2.0	n/a
1	Opt B:R	0.58	n/a	1.4	n/a
	Opt B:R + NB	–	n/a	–	–
	Fixed	0.62	50th	1.0, 2.0	0.13
2	Opt B:R	0.67	56th	1.1, 1.8	0.095
	Opt B:R + NB	0.63	38th	1.1, 1.7	0.043
	Fixed	0.59	33rd, 67th	0.67, 1.3, 2.0	0.27
3	Opt B:R	0.70	28th, 69th	1.0, 1.4, 1.9	0.11
	Opt B:R + NB	0.64	15th 43rd	0.8, 1.3, 1.7	0.041
	Fixed	0.54	25th, 50th, 75th	0.5, 1.0, 1.5, 2.0	0.34
4	Opt B:R	0.71	15th, 53rd, 79th	0.9, 1.2, 1.6, 2.0	0.12
	Opt B:R + NB	0.66	8th, 23rd, 46th	0.9, 1.1, 1.4, 1.7	0.037
	Fixed	0.72	n/a	n/a	n/a
FI	Opt B:R	–	–	–	–
	Opt B:R + NB	–	–	–	–

PTA = probability of target attainment (a $c_{ss,ave}$ between 1 and 2 *mg*/L); cut-score = the percentile of the distributions of *z* for the dose band cut-score; PrNB = the probability of non-beneficent outcomes (iatrogenic therapeutic failure caused by lack of dose intensity); Opt B:R = optimisation of dose banding based on net benefit assessment (as per Equation 7); Opt B:R + NB = optimisation of dose banding based on the composite of net benefit and non-beneficence assessment (as per Equation 9); FI = fully individualised.

[a]PTA for Fixed and Opt B:R is given by PTA_1 (Equation 7), otherwise when including NB is given by PTA_2 (Equation 9).

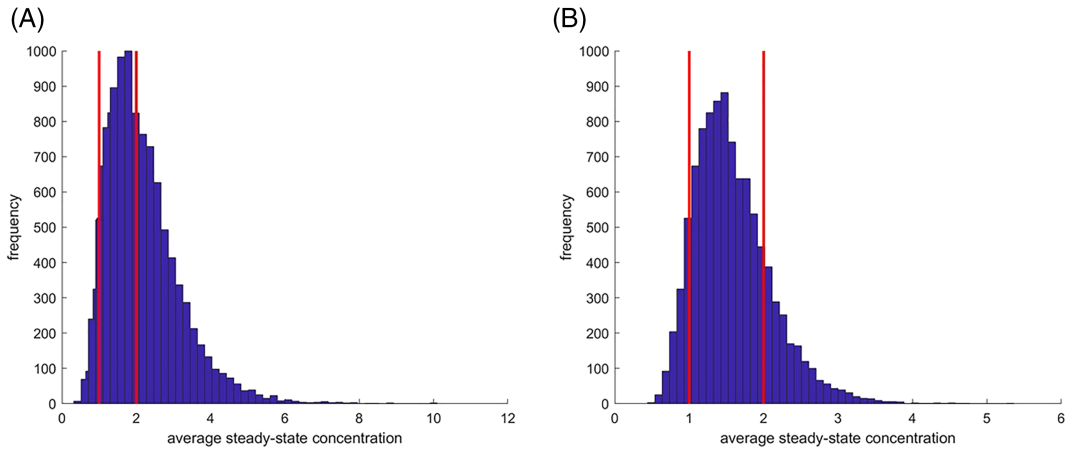

FIGURE 2 The frequency distribution of $c_{ss,ave}$ for (A) a single non-optimised dose level (one-dose-fits-all at a dose of 2 mg) and (B) covariate-based dosing (full dose individualisation). The vertical red lines represent the therapeutic range

virtual patients are shown in Figure 2a for the one-dose-fits-all scenario and Figure 2b for the fully individualised dosing. Optimising for net benefit only improved the PTA compared to fixed dose banding for all cut-scores.

When optimising for maximising the composite of both net benefit and minimisation of non-beneficence, we see a slight attenuation in the PTA compared to optimising for net benefit only but a significant reduction in the probability of non-beneficent outcomes (see Figure 3).

There were negligible improvements when increasing the number of dose levels, when considering net benefits only, beyond three dose levels and for the composite criteria after two dose levels (Figure 3). This was particularly evident for the composite criteria (that considered non-beneficence) where a constriction of the percentiles of the cut-scores (the points in the distribution of the covariate z where a dose change would be made) was seen with increasing dose levels, and adding more dose bands did not extend the range of doses. It was seen that the cut-scores for more than two dose levels were below the 50th percentile of z. When the number of dose bands extended beyond three, the number of unique doses remained very similar with dose levels within ±10% (essentially equivalent to a dose band being repeated). In contrast, the standard net benefit optimisation resulted in a wider range of doses and wider range of cut-scores as the number of dose levels increased. This is expected since the larger the range of doses, the more each patient can be individualised.

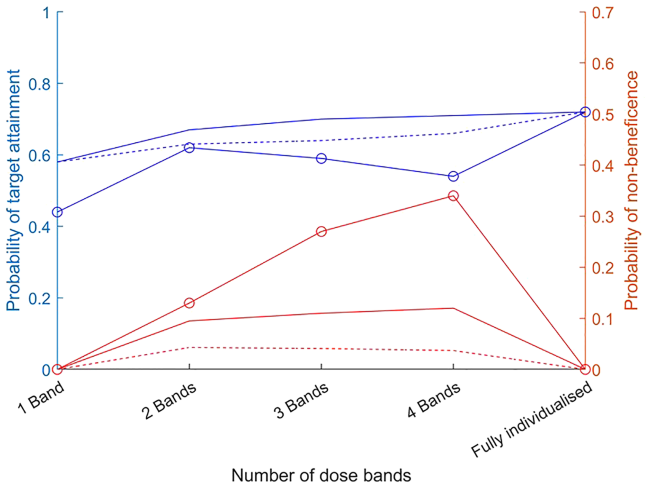

FIGURE 3 The probability of target attainment (the proportion of $c_{ss,ave}$ within the range of 1 – 2 mg/L) is shown on the left axis for: the dose banding optimised for net benefit (solid blue line), optimised for the composite of net benefit and minimising non-beneficence (dashed blue line) and fixed dosing based on principles of dose proportionality (open circle blue line). The probability of non-beneficent outcome (the proportion of $c_{ss,ave}$ that fall below the lower bound (1 mg/L) due to a reduction in dose intensification caused by dose banding) is shown in the right axis for: optimised for net benefit (solid red line), the composite criteria of net benefit and non-beneficence (dashed red line) and not optimised (open circle red line). The composite criteria represent a non-dominated solution

Importantly, the fixed dose bands, based on reducing the dose proportionately to the value of the covariate, provided poorer attainment of success and considerably higher probability of non-beneficent outcomes than an optimised strategy. This was particularly evident at the greater number of cut-scores.

4 | DISCUSSION

In this work various strategies were considered for individualising the starting dose, in terms of their utility to achieve a concentration within the desirable concentration range (i.e., where the $c_{ss,ave}$ fell within the therapeutic range). These included one-dose-fits-all (the dose is not individualised), dose banding, through to fully individualised dosing based on the patient's covariates. In either of these extreme cases (one dose fits all or fully individualised dosing), there is no opportunity loss due to non-beneficence since either no dose adjustments were made or the dose adjustment was entirely individualised to the patient. In contrast, dose banding, where patients are allocated to a particular dose level, resulted in some patients achieving success while others risked therapeutic failure. When the number of bands was increased, the probability of target attainment also increased for both optimised dose band settings but not for the fixed dose setting. Three methods of calculating the cut-scores for dose banding and the dosing level for each band were considered: (1) empirical dose banding, (2) dose banding optimised for therapeutic benefits and (3) dose banding optimised for a composite of both therapeutic benefits and non-beneficence. Empirical dose banding was based on standard PK principles, where.

$$CL \propto z, \text{ and } Dose \propto CL.$$

therefore

$$Dose \propto z.$$

Hence if the value of z was half of the normal value, then the value of dose rate would be halved accordingly. This is a standard approach for empirically dose adjusting (e.g., for renal function[7]) and is commonly applied to dose banding. We see here that this approach, when applied as fully individualised dosing, performed well (e.g., a mg/kg dose), but, when applied to dose banding, performed worse than optimised dose banding for any given number of bands in terms of both probability of target attainment and probability of non-beneficence. As an analogy, if the covariate z were creatinine clearance and assuming that normal creatinine clearance was 120 mL/min, then according to Table 3 the fixed cut-scores for dose change if we had two cut points would be creatinine clearance values of 40 mL/min and 80 mL/min. The fixed cut-scores would perform worse than optimised dose bands and equivalent to an optimised one-dose-fits-all setting. In addition, the fixed cut-scores would carry a 27% chance of iatrogenic failure (non-beneficence) by recommending a dose reduction when it was not needed. Continuing with this analogy, optimising the cut-

scores for net benefit (only) did not significantly alter the cut-scores but resulted in an increase in the dose level for the lower dose. Optimising for the composite of net benefit and non-beneficence also resulted in an increase in the lower dose level but compressed the cut-scores so that everyone over the 50th percentile (60 mL/min creatinine clearance) would receive the same dose. This is similar to the current practice with metformin dosing (as per the drug label).

When dose banding was extended to three cut-scores (four dose levels), we see a wider range of cut-scores for those optimised for net benefit (15th to 79th percentiles of the covariate) while maintaining the dosing range from 0.9 to 2.0 mg (only marginally wider than the two cut-score result). In contrast when optimising for the composite of net benefit and non-beneficence, the main driver appeared to be the lower cut-scores again with no cut-score above the 50th percentile. This is particularly evident when we compare the highest cut-score for the optimised composite at the 46th percentile compared to the net benefit only which was the 79th percentile. This suggests that a wide range of cut-scores are needed for net benefit whereas lower cut-scores, while maintaining dose intensity, improve non-beneficence.

Optimising for therapeutic benefit only (i.e., balancing benefits and risks according to Equation 7) both improved the probability of therapeutic benefit and also reduced the probability of non-beneficence compared to empirical dose banding. Optimising for the composite criterion of both therapeutic benefit and reducing non-beneficence provided a slightly attenuated probability of attaining therapeutic success but significantly reduced the probability of non-beneficence and represents a non-dominating solution and the best overall trade-off given our constraints. The trade-off is caused by the result that doing good for some patients will cause harm due to opportunity loss in others. A balance therefore needs to be sought between doing good and harming by not doing good. This concept has received little attention in the literature of dose banding. Importantly, however, covariate-based dosing was superior to all forms of dose banding and should be considered wherever practical and appropriate.

4.1 | Limitations and assumptions

In this work a very simple pharmacokinetic model was used which solved for the average steady-state concentration only and hence is based on only one parameter. This might be useful for considering a chronically dosed drug for which the area under the curve (AUC) is the primary driver for an exposure–response relationship of interest. This work is not intended to be extrapolated directly to a real-world setting, but does nevertheless show important implications. It is important to note that optimisation of dose bands is case-specific but the concepts described here can be generalised to other settings.

A simple therapeutic loss function was chosen as a square wave with boundaries defined by favourable net benefit assessment (i.e., one where the benefits outweighed the risks) which were based on the steady-state average concentration. More realistic loss functions, which may include skewness, could also be considered. The simplifying feature of the therapeutic range relates to the occurrence of side effects (maleficence). Here maleficence was only considered to occur when the $C_{ss,ave}$ exceeded the upper bound (2 mg/L). Obviously in real-world examples, side effects can occur at any exposure level but, for type B reactions—those that are idiosyncratic and considered not predictable—the probability increases with dose. This could be explored by applying this to an actual drug example in which the beneficent effects and maleficent effects were determined independently from models.

The covariate model was based on a linear causal relationship between clearance and the covariate (z). No misspecification of this relationship was considered nor was any error in the measurement of z considered. There are few (if any) truly linear causal relationships (e.g., CL and weight is nonlinear[1]), with the exception of perhaps glomerular filtration rate (GFR) on renal clearance for non-secreted drugs. However, this relationship is confounded by our approximation of GFR based on creatinine clearance equations and that creatinine is eliminated by both filtration and secretion mechanisms.[17] Nevertheless, the initial assumption of linearity provides an effective starting point and misspecification of empirical covariate relationships are not explored widely in current population PK models where relationships are generally based on and quantified by clinical data. The full covariate model included an unexplained between-subject variability on CL of 30% which is reasonable given the work of al-Sallami et al.[18] The simplifying assumption of no measurement error in z eliminates the possibility of an incorrect decision to change dose intensity. Although this is potentially important at values of z close to the cut-score, this was not explored as a simulation scenario.

Finally, we did not consider the resolution of the dose formulations. All formulations, whether injectable or solid dose, will have a pre-defined granularity in viable dose sizes. It is likely therefore that the benefits of dose individualisation by full individualisation based on covariates or dose banding will be lower in real-word data than in the current work. In this regard, post-marketing dose banding will be less effective than dose banding that informed dosage formulation, highlighting the importance of model-informed drug development.

5 | CONCLUSION

The influence of considering non-beneficence in choosing the dosing level in dose banding is explored in this work. This is likely to be of importance when few bands are determined (e.g., only one or two cut-scores). There is a need to explore this concept in future analyses of real-world data with more complicated pharmacokinetic-pharmacodynamic models.

ACKNOWLEDGEMENTS

The author is grateful for discussions with Prof. Mats Karlsson (Uppsala University) and feedback on the manuscript and discussions with Assoc. Prof. Rob Bies (University of Buffalo).

No funding was received for this work. Open access publishing facilitated by University of Otago, as part of the Wiley - University of Otago agreement via the Council of Australian University Librarians.

COMPETING INTERESTS

The author declares no conflict of interest.

DATA AVAILABILITY STATEMENT

No research data were used in this study. All research data were simulated.

ORCID

Stephen Duffull https://orcid.org/0000-0002-6545-9408

REFERENCES

1. Svensson EM, Yngman G, Denti P, McIlleron H, Kjellsson MC, Karlsson MO. Evidence-based design of fixed-dose combinations: principles and application to pediatric antituberculosis therapy. *Clin Pharmacokinet* 2018;57(5):591–599. doi:10.1007/s40262-017-0577-6
2. Lesko LJ, Schmidt S. Individualization of drug therapy: history, present state, and opportunities for the future. *Clin Pharmacol Ther* 2012;92(4):458–466. doi:10.1038/clpt.2012.113
3. Ogungbenro K, Patel A, Duncombe R, Nuttall R, Clark J, Lorigan P. Dose rationalization of pembrolizumab and nivolumab using pharmacokinetic modeling and simulation and cost analysis. *Clin Pharmacol Ther* 2018;103(4):582–590. doi:10.1002/cpt.875
4. Rothwell PM, Cook NR, Gaziano JM, Price JF, Belch JFF, Roncaglioni MC, Morimoto T, Mehta Z Effects of aspirin on risks of vascular events and cancer according to bodyweight and dose: analysis of individual patient data from randomised trials. *Lancet* 2018;392(10145):387–399. doi:10.1016/S0140-6736(18)31133-4
5. Jönsson S, Karlsson MO. A rational approach for selection of optimal covariate-based dosing strategies. *Clin Pharmacol Ther* 2003;73(1):7–19. doi:10.1067/mcp.2003.2
6. Viberg A, Cars O, Karlsson MO, Jönsson S. Estimation of cefuroxime dosage using pharmacodynamic targets, MIC distributions, and minimization of a risk function. *J Clin Pharmacol* 2008;48(11):1270–1281. doi:10.1177/0091270008320923
7. Doogue MP, Polasek TM. Drug dosing in renal disease. *Clin Biochem Rev*. 2012;32(2):69-73.
8. Larsson A, Akerstedt T, Hansson LO, Axelsson J. Circadian variability of cystatin C, creatinine, and glomerular filtration rate (GFR) in healthy men during normal sleep and after an acute shift of sleep. *Chronobiol Int* 2008;25(6):1047–1061. doi:10.1080/07420520802553614
9. Joffe M, Hsu CY, Feldman HI, Weir M, Landis JR, Hamm LL, Chronic Renal Insufficiency Cohort (CRIC) Study Group. Variability of creatinine measurements in clinical laboratories: results from the CRIC study. *Am J Nephrol* 2010; 31(5), 426–434. doi:10.1159/000296250
10. Brunetti L, Back H, Yu S, Jalil U, Kagan L. Evaluation and enhancement of standard equations for renal function estimation in individuals with components of metabolic disease. *BMC Nephrol* 2021;22(1):389. doi:10.1186/s12882-021-02588-4
11. Kuan I, Wilson LC, Leishman JC, et al. Metformin doses to ensure efficacy and safety in patients with reduced kidney function. *PLoS ONE* 2021;16(2):e0246247. doi:10.1371/journal.pone.0246247
12. Wright DF, Duffull SB, Merriman TR, Dalbeth N, Barclay ML, Stamp LK. Predicting allopurinol response in patients with gout. *Br J Clin Pharmacol* 2016;81(2):277–289. doi:10.1111/bcp.12799
13. Venkataramani AS, Bor J, Jena, AB. Regression discontinuity designs in healthcare research. *BMJ (Clinical Research Ed)* 2016;352:i1216. doi:10.1136/bmj.i1216
14. Beauchamp TL, Childress JF. *Principles of biomedical ethics*. New York: Oxford University Press; 2009.
15. Habibzadeh F, Habibzadeh P, Yadollahie M. On determining the most appropriate test cut-off value: the case of tests with continuous results. *Biochem Med (Zagreb)* 2016; 26(3):297–307. doi:10.11613/BM.2016.034
16. Holford NH, Buclin T. Safe and effective variability—a criterion for dose individualization. *Ther Drug Monit* 2012;34(5):565–568. doi:10.1097/FTD.0b013e31826aabc3
17. Ciarimboli G, Lancaster CS, Schlatter E, Franke RM, Sprowl JA, Pavenstädt H, Massmann V, Guckel D, Mathijssen RH, Yang W, Pui CH, Relling MV, Herrmann E, Sparreboom A. Proximal tubular secretion of creatinine by organic cation transporter OCT2 in cancer patients. *Clin Cancer Res* 2012;18(4):1101–1108. doi:10.1158/1078-0432.CCR-11-2503
18. al-Sallami HS, Cheah SL, Han SY, Liew J, Lim J, Ng MA, Solanki H, Soo RJ, Tan V, Duffull SB Between-subject variability: should high be the new normal? *Eur J Clin Pharmacol* 2014;70(11):1403–1404. doi:10.1007/s00228-014-1740-8

How to cite this article: Duffull S. Dose banding – Weighing up benefits, risks and therapeutic failure. *Br J Clin Pharmacol*. 2022;88(7):3474-3482. doi:10.1111/bcp.15307