Check for updates

Received: 13 October 2022 | Revised: 28 June 2023 | Accepted: 30 July 2023

DOI: 10.1002/bimj.202200285

Biometrical Journal

DISCUSSION

Reflections on the concept of optimality of single decision point treatment regimes

Trung Dung Tran[1] | **Ariel Alonso Abad**[2] | **Geert Verbeke**[2,3] | **Geert Molenberghs**[2,3] | **Iven Van Mechelen**[1]

[1]Quantitative Psychology and Individual Differences, KU Leuven, Leuven, Belgium

[2]I-BioStat, KU Leuven, Leuven, Belgium

[3]I-BioStat, Universiteit Hasselt, Hasselt, Belgium

Correspondence
Iven Van Mechelen, Quantitative Psychology and Individual Differences, KU Leuven, B-3000 Leuven, Belgium.
Email: Iven.VanMechelen@kuleuven.be

Funding information
Research Foundation - Flanders, Grant/Award Number: G080219N; Research Fund KU Leuven, Grant/Award Number: C14/19/054



This article has earned an open data badge "**Reproducible Research**" for making publicly available the code necessary to reproduce the reported results. The results reported in this article could fully be reproduced.

Abstract

In many areas, applied researchers as well as practitioners have to choose between different solutions for a problem at hand; this calls for optimal decision rules to settle the choices involved. As a key example, one may think of the search for optimal treatment regimes (OTRs) in clinical research, that specify which treatment alternative should be administered to each patient under study. Motivated by the fact that the concept of optimality of decision rules in general and treatment regimes in particular has received so far relatively little attention and discussion, we will present a number of reflections on it, starting from the basics of any optimization problem. Specifically, we will analyze the search space and the to be optimized criterion function underlying the search of single decision point OTRs, along with the many choice aspects that show up in their specification. Special attention is paid to formal characteristics and properties as well as to substantive concerns and hypotheses that may guide these choices. We illustrate with a few empirical examples taken from the literature. Finally, we discuss how the presented reflections may help sharpen statistical thinking about optimality of decision rules for treatment assignment and to facilitate the dialogue between the statistical consultant and the applied researcher in search of an OTR.

KEYWORDS

criterion function, optimal treatment regime, precision medicine, search space

1 | INTRODUCTION

In numerous applied domains, often multiple solutions are available for a particular problem. For example, in the clinical domain, to treat patients suffering from breast cancer, different therapies such as surgery, chemotherapy, radiation therapy can be administered (Kurian et al., 2014). Another clinical example occurs when first-line drugs for patients with hypertension can be prescribed by clinicians from among thiazides, beta-blockers, calcium channel blockers, angiotensin II receptor blockers, alpha-blockers, and so forth (Wright et al., 2018). In the educational domain, to teach students about cardiopulmonary resuscitation skills, different instructional methods such as traditional instruction, a case-based, or a web-based method can be used (Saraç & Ok, 2010). Finally, in the marketing domain, different strategies can be adopted to increase customer retention, including price incentives (such as an offer to put the customer on a more economical plan) and nonprice incentives (such as product improvement) (Ascarza et al., 2018).

In these and similar cases, there is a clear need for decision rules that indicate which treatment alternative (instructional method, retention strategy, respectively) should be administered to each patient (student, customer, respectively) under

 © 2023 Wiley-VCH GmbH.

study. Such rules can originate from different bases, including substantive hypotheses or theories (e.g., pharmacokinetic, pathophysiological, learning, or consumer behavior theories) and empirical data (from, e.g., randomized controlled trials or observational studies). It should be noted that, when empirical data are used to induce decision rules or to test hypotheses on them, the applied researcher or practitioner may wish to rely on a collaboration with a statistician or statistical consultant.

A fundamental concern when considering such decision rules is that they should be *optimal* from some point of view. Importantly, the viewpoint in question should reflect what the clinician, educational scientist, or marketeer cares about. In other words, a decision rule should meet the specific substantive concerns involved. For example, if in the hypertension example mentioned above, the clinician is concerned about the proportion of patients for whom blood pressure falls below 140/90 after taking the drug, optimality should be driven by that proportion. Similarly, if in the marketing example the company cares about the overall net financial consequences as implied by the increased customer retention minus the cost of the retention strategy, optimality should be directed by these overall financial consequences.

In line with the above, the issue of optimal decision rules is relevant in many disciplines. As will be further clarified in this paper, a common conceptual basis can be used to handle this issue across all domains under consideration. To simplify our presentation, however, we will further put our primary substantive focus on medicine and biostatistics. Hence, we will use primarily terminology from clinical biostatistics, such as patients, treatments, and treatment alternatives. In addition, throughout the text we will use the biostatistical term "optimal treatment regimes (OTR's)" to denote optimal decision rules. That being said, the translation of the conceptual basis outlined in this paper to other domains will be straightforward.

There is an extensive literature in biostatistics on OTRs. The interested reader may be referred to several recent books (Chakraborty & Moodie, 2013; Kosorok & Moodie, 2015; Tsiatis et al., 2020) and review articles (Chakraborty & Murphy, 2014; Laber, Lizotte, Qian, et al., 2014; Schulte et al., 2014). It should be noted, however, that a large part of the literature on OTRs focuses on rather complex and technical aspects such as the estimation of OTRs based on a single optimality criterion. At this point, most often the optimality criterion at hand is considered as a given without much attention or discussion being spent at it. In addition, whereas a number of different possible specifications of the concept of treatment regime optimality have been proposed, in general their appearance is rather scattered. Hence, the current status of literature may cause difficulties for the applied researcher or statistical consultant who wish to select an appropriate specification of optimality when facing particular substantive concerns.

To address this problem, in this paper the concept of treatment regime optimality will be given the attention that it properly deserves. Specifically, we will present reflections on treatment regime optimality, starting from the basics of the broad category of optimization problems that are ubiquitous in statistics, data analysis, biometrics, and related disciplines. In this way, we aim at (i) sharpening statistical thinking about optimality of decision rules for treatment assignment and (ii) linking optimality criteria to substantive concerns. As such we hope to facilitate the dialogue during the conceptual and planning stages between statisticians or statistical consultants on the one hand and applied researchers or practitioners in search of an OTR on the other hand, with all people involved in such a dialogue constituting the target audience for this paper. Note that we do not envision statisticians in search of technical details on a precision medicine analysis approach, but statisticians and a general audience who want to reflect on the concept of optimality of treatment regimes. The focus of the paper will therefore be conceptual in nature, whereas important estimation and implementation aspects of a precision medicine analysis approach will be left aside. To safeguard this focus, we will further strictly limit the discussion in our paper to decision rules on treatment choices that involve a single decision point only, while leaving more complex setups that involve multiple decision points (i.e., dynamic treatment regimes) aside.

The remainder of this paper is organized as follows: In Section 2, we will present a conceptual basis for treatment regime optimality, starting from the fundamentals of optimization problems in general, while also taking into account the link with substantive concerns. In Section 3, we will illustrate with three instantiations of the conceptual basis taken from literature. The paper will conclude with some discussions in Section 4 and a final remark in Section 5.

2 | TREATMENT REGIMES: ANALYSIS OF THE CONCEPT OF OPTIMALITY

2.1 | Preliminaries and notation

Let $\mathcal{T}$ be a target set of patients of interest in a particular study, where an arbitrary patient is denoted by p. For example, we can consider the target set of patients suffering from hypertension or the target set of patients suffering from breast cancer.

15214036, 2023, 8, Downloaded from https://onlinelibrary.wiley.com/doi/10.1002/bimj.202200285 by <Shibboleth>-member@kuleuven.be, Wiley Online Library on [08/01/2024]. See the Terms and Conditions (https://onlinelibrary.wiley.com/terms-and-conditions) on Wiley Online Library for rules of use; OA articles are governed by the applicable Creative Commons License

Further, among the patients with breast cancer, one may optionally focus on the target set of patients being oestrogen receptor positive, that is to say, patients whose cancer cells have receptors for oestrogen (Nicholson et al., 1981).

Let $\mathcal{A}$ further denote the set of all treatment alternatives that are to be considered at the decision point in question (which optionally may include the alternative of doing nothing, or assignment to a placebo). Obviously, $\mathcal{A}$ should contain at least two treatment alternatives, otherwise there is no decision issue. For instance, in the breast cancer example, $\mathcal{A}$ may comprise surgery, chemotherapy, radiation therapy, and so forth. Note that $\mathcal{A}$ can comprise a finite or an infinite number of treatment alternatives, with continuous dose of some medication being an obvious example of an infinite number of treatment alternatives (Demonty et al., 2009). When $\mathcal{A}$ is finite, numbers are often used to denote the treatment alternatives (e.g., $\mathcal{A} = \{0, 1\}$ when there are two treatment alternatives involved only).

We may now formally define a treatment regime as a decision rule that links every patient in the target set of patients to a particular treatment alternative. Mathematically, given a target set $\mathcal{T}$ of patients and some set of treatment alternatives $\mathcal{A}$, a treatment regime r is defined as a function from $\mathcal{T}$ to $\mathcal{A}$:

$$\begin{aligned} r : \mathcal{T} &\to \mathcal{A} \\ p &\mapsto a \end{aligned}.$$

The set of all treatment regimes starting from $\mathcal{T}$ is further denoted by $\mathcal{R}_{\mathcal{T}\to\mathcal{A}}$.

Obviously, as will be further discussed below, treatment regimes will very often factor in baseline patient-related information that is available at the point of deciding on a treatment alternative. This may include both personal characteristics of the patient (such as gender, age, specific type of disease, and genetic information) and characteristics of the patient's family or, more broadly, the patient's environment. All available patient-related information will further be referred to as covariates and will be denoted by $\mathbf{X}$, with $\mathbf{X}(p)$ denoting the covariate pattern of patient p.

A final essential key ingredient to compare different treatment regimes and to define treatment regime optimality is information on the outcome after a patient underwent one of the different treatment alternatives. Whereas, in general, interest may be in multiple outcome variables (which may capture several beneficial consequences of the treatment as well as several undesirable side effects), for simplicity's sake, the primary focus of this paper will be on setups that involve a single quantitative outcome variable only. That being said, let Y denote the quantitative outcome variable of interest, where higher values of Y are assumed to be more beneficial (without loss of generality). Let $Y_a^*(p)$ further denote the potential outcome that we would observe if subject p were to receive treatment a (Neyman, 1923; Rubin, 1974, 1978, 2005), where we assume that this outcome is a fixed number that is not subject to within-person variability (which may be considered an oversimplification: see Senn, 2018). Note that the asterisk and the specification "potential" reflect the fact that, in many typical clinical settings, for a particular patient we can observe at most one outcome, as that patient can receive at most one treatment alternative.

Finally, we may define the potential outcome under treatment regime r, which will further be denoted by Y_r^*. Specifically, if r assigns treatment alternative a to patient p (i.e., $r(p) = a$), then $Y_r^*(p)$ is defined as $Y_a^*(p)$. In other words, the random variable Y_r^* on $\mathcal{T}$ is defined as follows:

$$Y_r^*(p) = Y_{r(p)}^*(p).$$

This means that the potential outcome under treatment regime r for patient p is the potential outcome of that patient under the treatment alternative prescribed to that patient by r.

2.2 | An optimization problem

Broadly speaking, searching for an OTR is to look, within a set of treatment regimes under study, for the regime that optimizes some specific objective function, such as to maximize some kind of overall benefit or profit. For example, in the hypertension example, one might wish to search within a prespecified set of treatment regimes for the treatment regime that reduces as much as possible the patients' blood pressure to a normal level. For breast cancer patients, the objective might be to increase as much as possible the survival time. Thus, there are two key components that are to be considered when searching for OTRs: the to be prespecified set of treatment regimes within which we look for optimal ones, and the objective function that is to be optimized.

15214036, 2023, 8, Downloaded from https://onlinelibrary.wiley.com/doi/10.1002/bimj.202200285 by <Shibboleth>-member@kuleuven.be, Wiley Online Library on [08/01/2024]. See the Terms and Conditions (https://onlinelibrary.wiley.com/terms-and-conditions) on Wiley Online Library for rules of use; OA articles are governed by the applicable Creative Commons License

Formally speaking, looking for an OTR then is an instance of a generic search for the solution of a mathematical optimization problem where we look within a search space for the element that maximizes some criterion or objective function. Specifically, given some target set of patients $\mathcal{T}$, an OTR, denoted by r^{opt}, is a treatment regime that satisfies:

$$r^{opt} = \underset{r \in \mathcal{R}_s}{\operatorname{argmax}}\, C(Y_r^*), \tag{1}$$

where $\mathcal{R}_s$, a to be prespecified subset of $\mathcal{R}_{\mathcal{T}\to\mathcal{A}}$, is the search space, and C is the to be prespecified criterion or objective function, where, without loss of generality, higher values of C are assumed to be more preferable. Otherwise, whereas one might read (1) to suggest that there is a single OTR that maximizes objective function $C(Y_r^*)$, in general there is a possibility of ties, with several treatment regimes yielding the same maximal value of the function in question.

We will carefully discuss the choices of $\mathcal{R}_s$ and C in Sections 2.3 and 2.4, respectively. It is important to emphasize already now that these choices should primarily be guided by the particular interests and concerns of the applied researcher or clinician.

2.3 | The search space

To specify a search space, a target set and a set of treatment regimes $\mathcal{R}_s$ need to be chosen. The specification of a well-defined target population of patients, which may require a sound substantive-theoretical justification, is a standard issue in biostatistics. The choice of a set of treatment regimes is somewhat more involved.

The simplest type of regime is the one-size-fits-all one, meaning that the same treatment alternative is to be assigned to all patients in $\mathcal{T}$. Denote by $R_{T\to\{a\}}$ the set of all one-size-fits-all treatment regimes, with

$$R_{T\to\{a\}} = \{r : \mathcal{T} \to \{a\}, \text{ for some } a \in \mathcal{A}\}. \tag{2}$$

A one-size-fits-all approach is preferable if one and the same treatment alternative outperforms all other alternatives for every individual patient. There are, however, clear indications that this may not always be the case (Doestzada et al., 2018; Flashner-Abramson et al., 2019; Kiely, 2015). This motivates the development of so-called "personalized" treatment regimes, which allow for a differentiation between patients in treatment assignment, meaning that different patients can be assigned to different treatment alternatives.

Clearly, the ultimate ideal of a personalized approach is to assign each patient to the treatment alternative that yields the best possible (potential) outcome for that patient (Bose & Freedman, 2020; Trivedi, 2016; Whitcomb, 2019). This implies that, in general, a fully fledged personalized approach requires to consider the full set of treatment regimes $\mathcal{R}_{\mathcal{T}\to\mathcal{A}}$. However, there is a limitation on this in practice. Indeed, at the moment of deciding on a treatment alternative, typically none of the potential outcomes is known for the patient at hand. This implies that treatment assignment decisions can be based only on patient-related covariate information that is available at that moment. This further implies that two patients with the same pattern of covariate information should necessarily be assigned to the same treatment alternative; therefore, the level of the unique individual patient remains out of reach (which is why the US National Research Council, 2011, expressly recommends to replace the term "personalized medicine" by the term "precision medicine"). In other words, instead of working with the full set of regimes $\mathcal{R}_{\mathcal{T}\to\mathcal{A}}$, in practice we may only consider regimes that can be factored via the range of all covariate patterns, $\mathbf{X}(\mathcal{T})$, that is to say, the family of regimes $\mathcal{R}_{\mathcal{T}\to\mathbf{X}(\mathcal{T})\to\mathcal{A}}$.

Next, the function from $\mathbf{X}(\mathcal{T})$ to $\mathcal{A}$ needs to be specified. A first option to do so is to start from a *prespecified* partition of $\mathbf{X}(\mathcal{T})$ and to subsequently look for the best treatment alternative for each resulting partition class. Alternatively, one may estimate the function that yields an optimal treatment regime, which will induce a partitioning of $\mathbf{X}(\mathcal{T})$ (and $\mathcal{T}$) into post hoc subgroups. For this estimation, one may note that, except in a limited number of cases that involve a few discrete covariates only, each of which can only take a very small number of values, the set of functions from $\mathbf{X}(\mathcal{T}) \to \mathcal{A}$ will be very large in size. Hence, it will usually not be feasible to look for an optimal function within this set as a whole. As a way out, in practice, one may usually wish to limit the search to a prespecified class of functions $\mathcal{F}$ with a well-defined functional form. Examples include the class of linear functions (Brinkley et al., 2010; Qian & Murphy, 2011; Wang et al., 2018; Xu et al., 2015; Zhang, Tsiatis, Laber, et al., 2012; Zhang et al., 2013), the class of tree-based functions (Dusseldorp & Van Mechelen, 2014; Foster et al., 2011; Laber & Zhao, 2015; Sies & Van Mechelen, 2017; Zhang, Tsiatis, Davidian, et al.,

15214036, 2023, 8, Downloaded from https://onlinelibrary.wiley.com/doi/10.1002/bimj.202200285 by <Shibboleth>-member@kuleuven.be, Wiley Online Library on [08/01/2024]. See the Terms and Conditions (https://onlinelibrary.wiley.com/terms-and-conditions) on Wiley Online Library for rules of use; OA articles are governed by the applicable Creative Commons License

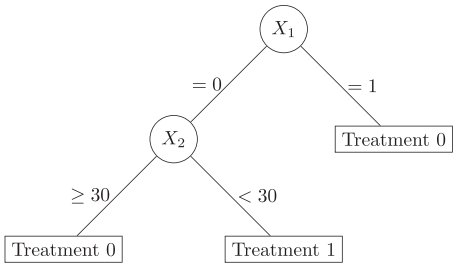

FIGURE 1 A hypothetical example of a tree-based treatment regime where X_1 is binary and X_2 is continuous.

2012), decision lists (Zhang et al., 2015), and the class of nonlinear or generalized linear classifiers (Imai & Ratkovic, 2013; Zhao et al., 2012; Zhang, Tsiatis, Davidian, et al., 2012).

Let us denote by $\mathcal{R}_{\mathcal{T}\to\mathbf{X}(\mathcal{T})\overset{\mathcal{F}}{\to}\mathcal{A}}$ the set of treatment regimes that is obtained by constraining the function from $\mathbf{X}(\mathcal{T})$ to $\mathcal{A}$ to belong to the prespecified family of functions $\mathcal{F}$. One may note that $\mathcal{R}_{\mathcal{T}\to\mathbf{X}(\mathcal{T})\overset{\mathcal{F}}{\to}\mathcal{A}} \subset \mathcal{R}_{\mathcal{T}\to\mathbf{X}(\mathcal{T})\to\mathcal{A}} \subset \mathcal{R}_{\mathcal{T}\to\mathcal{A}}$. Otherwise, if the family of functions $\mathcal{F}$ contains the constant function (which is true for most families), it further also holds that $\mathcal{R}_{\mathcal{T}\to\{a\}} \subset \mathcal{R}_{\mathcal{T}\to\mathbf{X}(\mathcal{T})\overset{\mathcal{F}}{\to}\mathcal{A}}$.

We illustrate by means of two common restricted classes of functions, namely, the linear and the tree-based ones. For simplicity's sake, we will do so for the case where $\mathcal{A}$ contains two treatment alternatives only, $\mathcal{A} = \{0, 1\}$. A linear treatment regime then takes the form $I(\beta_0 + \beta_1 X_1 + \cdots + \beta_k X_k > 0)$, with $I(.)$ denoting the indicator function, and with its argument and the negation of it defining the two post hoc subgroups. From its part, a tree-based treatment regime relies on a tree as a basis of a treatment assignment rule. A hypothetical example of such a tree is given in Figure 1, with the internal nodes of the tree being represented by circles (each of which contains a covariate), and with the terminal nodes or leaves of the tree being represented by rectangles (each of which contains a treatment alternative); in this case, the leaves constitute the post hoc subgroups.

The above implies that, apart from the target set of patients, choices to be made in the specification of a search space include the choice between the set of the one-size-fits-all and a set of precision regimes; in case precision regimes are chosen, subsequently a choice has to be made between regimes based on prespecified versus regimes based on post hoc subgroups of patients; if post hoc subgroups are chosen, further choices include the choice of a family of functions $\mathcal{F}$, which covariates have to be involved, and in which way they should be involved (e.g., with which transformations and in which functional form, such as, without or with covariate by covariate interaction terms). Possible grounds in these choices are a substantive-theoretical basis, simplicity (including ease of interpretation, estimation, and communication, and speed and cost of implementation in clinical practice), fit to the data, and replicability (with the latter linking up with the seventh statistical virtue advanced by Gelman & Hennig, 2017).

Specifically, with regard to the choice between one-size-fits-all versus precision regimes, one may argue that the former score better on simplicity (in all possible respects) and on stability/replicability, with (treatment by subgroup) interactions in general being harder to replicate than (treatment alternative) main effects. On the other hand, precision regimes may go with a better fit to the data, notably when treatment effect heterogeneity is in play, and in particular in presence of treatment–subgroup interactions that imply that which treatment alternative is optimal varies across subgroups of patients (Lipkovich et al., 2017).

With regard to the choice between precision regimes based on prespecified versus based on post hoc subgroups, one might argue that post hoc subgroups may lack a sound substantive-theoretical justification; otherwise, because of this, some authors (Feinstein, 1998; Rothwell, 2005) are strongly opposed to inducing post hoc subgroups from data, which they refer to as "data dredging." On the other hand, one could also argue that by allowing for post hoc subgroups, one may leave room for learning hitherto unknown insights from the data at hand.

As to the choice of a family of functions $\mathcal{F}$, one may note that families differ from one another with regard to: (a) simplicity (and, hence, ease of interpretation and communication), (b) stability/replicability, and (c) family flexibility/breadth (which is directly linked to fit to the data). For example, the linear family scores good on (b) and rather good on (a); the tree-based family further scores very good on (a) but less good on (b); and finally, families based on support vector regression models (Zhao et al., 2009), random forests (Wager & Athey, 2018), and Bayesian regression trees (Hahn et al., 2020) score very good on (c), but pretty low on (a).

15214036, 2023, 8, Downloaded from https://onlinelibrary.wiley.com/doi/10.1002/bimj.202200285 by <Shibboleth>-member@kuleuven.be, Wiley Online Library on [08/01/2024]. See the Terms and Conditions (https://onlinelibrary.wiley.com/terms-and-conditions) on Wiley Online Library for rules of use; OA articles are governed by the applicable Creative Commons License

Regarding which covariates have to be involved and how they should be so, higher numbers of covariates and allowing for more complex functional forms obviously imply larger search spaces; these go with a better fit to the data, but also with lower levels of parsimony and replicability. As an example, one may consider the search for an OTR in case of high-dimensional omics data (see, e.g., Gerstung et al., 2017), which may require the involvement of various kinds of covariate selection techniques (such as regularized regression methods).

We end this section with two concluding remarks. First, as discussed above, prior theoretical knowledge may be invoked at several points in the decision process, with some authors even dismissing decisions that are not fully based on it. At this point, one could draw a distinction between, on the one hand, a purely confirmatory analysis that is entirely based on prior knowledge and/or biologically plausible prior hypotheses, and, on the other hand, a purely exploratory analysis that does not factor in any such prior knowledge at all. Importantly, however, we believe that many intermediate cases could and should be considered, beyond such a confirmatory–exploratory dichotomy. Examples may include cases in which full prior theoretical knowledge or hypotheses are available on which covariates are relevant, in absence of precise hypotheses on the relevant types of preprocessing or functional forms. Also, one might consider analyses that involve a number of theoretically justifiable covariates, next to a few other covariates that may be included in the analyses for exploratory purposes only. Through such intermediate scenarios, one might wish to simultaneously test prior hypotheses and learn novel insights from the data.

A second remark pertains to replicability. Whereas, in general, a larger search space leads to a possibility of arriving at better regimes, it potentially also leads to a higher risk of a lack of replicability. This holds in particular if we move from a one-size-fits-all to a precision approach. In fact, nontrivial precision regimes implicitly rely on disordinal or crossed treatment alternative by patient subgroup interactions, with statistical inference about such interactions (and, hence, replicability of estimated precision regimes) being especially challenging (Brookes et al., 2004; Lee et al., 2015). This implies that one should take care to always subject results from an OTR exercise within a precision approach to well-designed replication studies before implementing them in practice (Friedman et al., 2015), still apart from the fact that replication studies are always needed, anyhow.

2.4 | Criterion functions

A second key component of the definition of OTRs in (1) is the objective or criterion function $C()$. As implied by (1), searching for an OTR comes down to comparing different candidate regimes r in terms of (marginal) aspects of the distribution of the corresponding random variables Y_r^* via criterion function $C()$. In this section, we will discuss how a suitable criterion function can be selected. For notational convenience, we will assume throughout, without loss of generality, that the target set is $\mathcal{T}$ and the search space $\mathcal{R}_s$ (a subset of $\mathcal{R}_{\mathcal{T}\to\mathcal{A}}$).

We start with a review of criterion functions that have been proposed in the literature. We will do so while emphasizing the distributional aspects they involve along with their substantive justification.

The most obvious option is to compare different treatment regimes "on average," with the OTR being the regime with the highest expected potential outcome (Brinkley et al., 2010; Chakraborty & Moodie, 2013; Qian & Murphy, 2011):

$$r^{opt} = \underset{r\in\mathcal{R}_s}{\operatorname{argmax}}\, E(Y_r^*). \tag{3}$$

Related to the above, clinicians sometimes want to evaluate therapy outcome in relation to a clinically meaningful threshold, which can be interpreted as the minimum threshold for success. For example, when evaluating the effectiveness of weight loss interventions, a 5% weight loss is sometimes considered as a minimum threshold for intervention success (Williamson et al., 2015). This naturally, leads to the following definition of an OTR (Imbens & Rubin, 2015; Jiang et al., 2017; Linn et al., 2017):

$$r^{opt} = \underset{r\in\mathcal{R}_s}{\operatorname{argmax}}\, P(Y_r^* \geq \tau), \tag{4}$$

where τ denotes the prespecified threshold. It should be noted that (4) is a special case of (3). Indeed, let Z be a discretized version of Y, with $Z = I(Y \geq \tau)$. It then immediately follows that $P(Y_r^* \geq \tau) = E(Z_r^*)$.

15214036, 2023, 8, Downloaded from https://onlinelibrary.wiley.com/doi/10.1002/bimj.202200285 by <Shibboleth>-member@kuleuven.be, Wiley Online Library on [08/01/2024]. See the Terms and Conditions (https://onlinelibrary.wiley.com/terms-and-conditions) on Wiley Online Library for rules of use; OA articles are governed by the applicable Creative Commons License

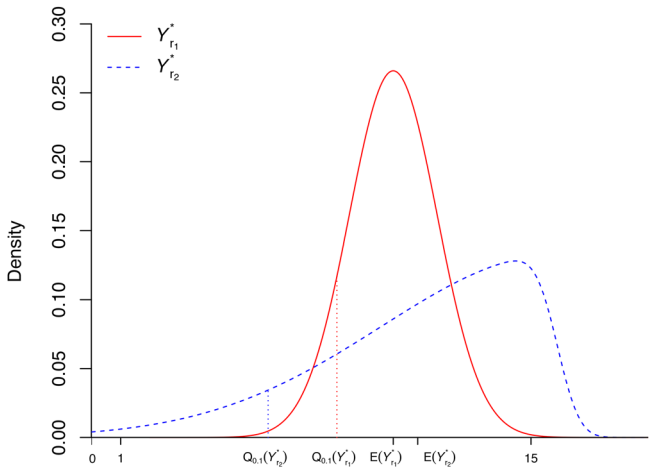

FIGURE 2 An artificial example depicting the distributions of the potential outcomes under two different treatment regimes.

A possible problem with the mean criterion is that it exclusively focuses on the average patient only, possibly at the expense of a large fraction of the target set. For example, whereas in Figure 2 treatment regime r_2 is on average better than treatment regime r_1, r_2 implies a substantially higher probability for a patient to end up in the left tail of the outcome distribution.

As a way out, one might consider to require a treatment regime to be beneficial for the target set as a whole; this could be achieved by maximizing the minimum of the distribution of Y_r^*. The latter, however, would imply an extreme and rather inconvenient choice, as the minimum of a distribution is hard to estimate. Instead of 100%, a more feasible solution could be to require the treatment regime to be as beneficial as possible for a large part P (such as, e.g., 90% or 95%) of the target set. Given some prespecified value of $q = 1 - P, 0 < q < 1$, this could be achieved by means of the following optimization problem (Imbens & Rubin, 2015; Wang et al., 2018):

$$r^{opt} = \underset{r \in \mathcal{R}_s}{\text{argmax}}\, Q_q(Y_r^*), \tag{5}$$

with $Q_q(.)$ denoting the qth quantile of the distribution under study. As an example, in Figure 2, r_1 outperforms r_2 in terms of the criterion function $Q_{.10}(.)$.

In still other substantive contexts, one might be concerned about variability in outcome across the target set. For instance, when devising redistribution or social transfer programs, one may be especially concerned about income spread, and in particular about the gap between rich and poor households (Caminada et al., 2012). Hence, this and similar concerns (also in a clinical context) may justify criterion functions based on (the inverse of) the variance in outcome. Obviously, most often it will be meaningless to achieve low variability in outcome by lowering the outcome for all members of the target set (e.g., to reduce social inequality by making all households poor). Hence, in such cases, one may also be concerned about the average outcome, which may motivate to consider criterion functions based on a combination of measures of location and spread (Zajonc, 2012) such as:

$$r^{opt} = \underset{r \in \mathcal{R}_s}{\text{argmax}}\, (\gamma_1 E(Y_r^*) - \gamma_2 SD(Y_r^*)), \tag{6}$$

where γ_1 and γ_2 are to be prespecified and positive constants (with $\gamma_1 + \gamma_2 = 1$) and where SD stands for standard deviation.

From the above review, one may derive a number of choices that are to be made in the specification of a criterion function:

(1) A first of these pertains to an optional preprocessing of the outcome variable Y. Examples include a dichotomization based on some threshold τ as in (4). As another example, one may refer to the calculation of the minimum of an observed survival time Y and some prespecified censoring limit τ (i.e., $\tilde{T} = \min(T, \tau)$) (with $\tilde{T}$ optionally further being built in a restricted mean survival time criterion function; Tsiatis et al., 2020). It should be noted that, within a

15214036, 2023, 8, Downloaded from https://onlinelibrary.wiley.com/doi/10.1002/bimj.202200285 by <Shibboleth>-member@kuleuven.be, Wiley Online Library on [08/01/2024]. See the Terms and Conditions (https://onlinelibrary.wiley.com/terms-and-conditions) on Wiley Online Library for rules of use; OA articles are governed by the applicable Creative Commons License

TABLE 1 Summaries of the distributions of the potential outcomes under the two treatment regimes depicted in Figure 2 (with optimal values printed in bold).

Regime	Mean	Prob. > 8	Median	10th quantile	SD	$0.7 \times \text{Mean} - 0.3 \times \text{SD}$
r_1	10.300	**0.937**	10.300	**8.378**	**1.500**	**6.760**
r_2	**11.136**	0.812	**11.853**	6.031	3.648	6.701

particular type of preprocessing of the response variable, possibly additional choices may be needed (such as, e.g., the choice of a threshold for success, which otherwise can even be made person-specific).

(2) A second and most critical choice aspect is that of the distributional characteristics one may wish to focus on in the specification of a criterion function. As illustrated above, characteristics such as measures of central tendency, quantiles, and measures of spread can be considered at this point, in addition to other distributional aspects such as skewness, with a distribution with a thick left tail probably being less desirable (see also Figure 2). Otherwise, within the choice of a particular characteristic, further choices may sometimes be needed, such as the value of the quantile of interest in (5) or the specific measure of variation opted for (e.g., the interquartile range or the 95th–5th percentile difference; Hanushek & Wäßmann, 2006; instead of the standard deviation in (6)).

(3) That being said, multiple distributional aspects can be simultaneously of interest. This is exemplified by (6). As a second example, if one cares about both the average performance and the lower tail of the distribution, a combination of mean and quantile criteria may be considered. Otherwise, it could also make sense to combine two instantiations of the same criterion type. As an example, one could be willing to look for regimes that yield simultaneously an as high as possible probability of exceeding some high threshold for success τ_1 and an as low as possible probability of falling below some low threshold for failure τ_2; this could be achieved by combining criterion function (4) twice, with two different thresholds $\tau_1 > \tau_2$:

$$\underset{r \in \mathcal{R}_s}{\operatorname{argmax}} \left(\gamma_1 P(Y_r^* \geq \tau_1) + \gamma_2 P(Y_r^* \geq \tau_2)\right). \tag{7}$$

Note further that, whenever a compound criterion function is chosen, a subsequent choice of the weights (e.g., γ_1 in (6) and (7), assuming that $\gamma_2 = 1 - \gamma_1$) is required.

Given the wealth of possible options for criterion functions, one may wonder whether the choice between them really matters, in that different choices may result in fairly different OTRs. The following artificial example illustrates that this is the case, indeed:

Example 2.1. In this example, we assume that $\mathcal{A} = \{1, 2\}$ and that a one-size-fits-all approach is of interest, implying that the search space equals $R_{T \to \{a\}} = \{r_1, r_2\}$. We further assume that the two potential outcomes for every individual in the target set under study are known (which of course will almost never be the case in practice). The distributions of the potential outcomes under the two treatment regimes are depicted in Figure 2; several summary characteristics of them are listed in Table 1 (with the code to calculate them being available in the Supporting Information, for reproducibility purposes). Using the mean criterion, r_2 is chosen. If, however, we want to maximize the probability for the outcome to exceed 8.0, r_1 is preferable, with a probability of 0.937 for patients to satisfy this criterion. Similar to the mean criterion, if the median criterion is used, r_2 is chosen. However, as can be seen in Figure 2, there is a large probability for patients to end up in the left tail under r_2. In fact, under r_1 there is a probability of 0.90 for an outcome to exceed 8.378, whereas the corresponding quantile for r_2 equals 6.031. Thus, if we care to help on as much as possible 90% of the target set under study, we should choose r_1. Finally, we can see that the outcomes under r_2 range from 0 to 18 while under r_1, they range just from 5 to 15 (which is also reflected by the corresponding standard deviations). Thus, although r_2 outperforms r_1 on average, it creates more inequality among the patients. Therefore, if we also take into account inequality, r_1 might be preferable (see rightmost column of Table 1).

We now turn to possible grounds for the different choices involved in the specification of a criterion function:

(1) Regarding a possible preprocessing of the outcome variable, in some cases a particular type of preprocessing may be mandatory. For example, in case of survival data, because of censoring, the mean criterion (3) cannot be readily applied

15214036, 2023, 8, Downloaded from https://onlinelibrary.wiley.com/doi/10.1002/bimj.202200285 by <Shibboleth>-member@kuleuven.be, Wiley Online Library on [08/01/2024]. See the Terms and Conditions (https://onlinelibrary.wiley.com/terms-and-conditions) on Wiley Online Library for rules of use; OA articles are governed by the applicable Creative Commons License

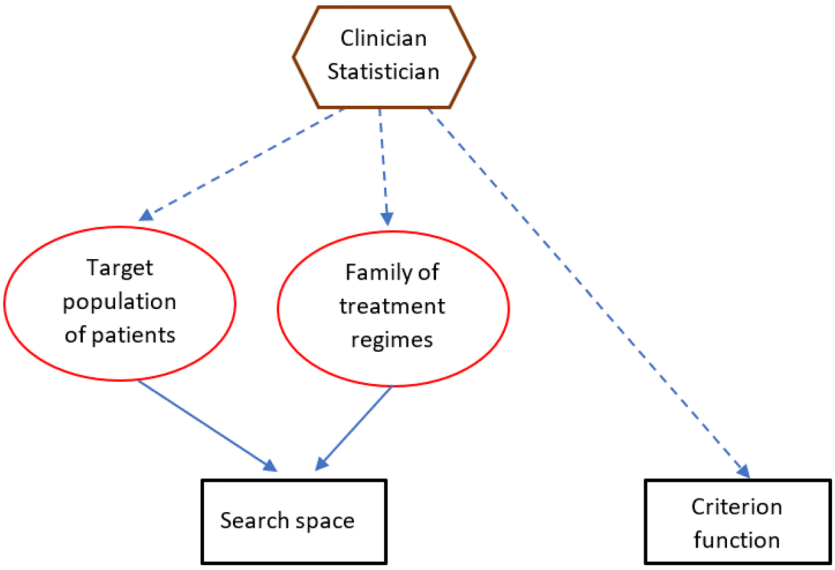

FIGURE 3 Aspects of treatment regime optimality to be thought and talked through in the conceptual/planning stages of optimal treatment regime (OTR) research (with different shapes for the actors, the two key ingredients of optimization, and the constituents of the first ingredient, viz., the search space).

to the raw data (Goldberg & Kosorok, 2012) and moving on to a restricted mean survival time is inevitable. In other cases, some kind of preprocessing may not be mandatory but desirable in view of substantive concerns. For instance, in some contexts with outcome variables that have been measured on scales with clinically meaningful marker values, the primary interest may be in the absolute rather than the relative magnitude of the values of the outcome. This may justify a preprocessing of the outcome based on thresholds corresponding to the marker values in question. As an example, one may think of systolic and/or diastolic mm Hg blood pressure as outcome variables in therapies for hypertension, which are largely agreed to be preprocessed based on absolute cutoff values such as 140 for systolic, and 90 for diastolic blood pressure.

(2) With regard to which distributional characteristics one may wish to focus on, this choice may typically be guided in part by the type of the outcome variable under study. For example, when the outcome is binary, the mean/probability criterion is a very natural (and perhaps the only meaningful) choice, whereas in case of continuous variables that may be subject to outliers, robust criteria such as (5) might be preferable over (3). More importantly, however, the choice of distributional characteristics should reflect the substantive concerns at stake. For example, if one cares about the left tail where patients with less desirable outcomes are located, a quantile criterion would be preferable. As a second example, one may think of educational programs, in which educationalists typically care about equality in access to academic resources as well as equality in pupil performance or knowledge acquisition. In all such cases, spread is of concern.

(3) Finally, the choice of a combination of distributional aspects will jump out whenever multiple substantive concerns simultaneously are at stake.

We conclude this section with a graphical summary representation (Figure 3) of the different aspects of treatment regime optimality that have to be thought and talked through by clinicians and statisticians in the conceptual and planning stages of OTR research (prior to the implementation and estimation stages).

3 | ILLUSTRATIVE EXAMPLES FROM LITERATURE

To illustrate the conceptual basis outlined above, we will revisit in this section three empirical examples taken from literature. For each example, after summarizing the main relevant points, we will discuss how it connects to our taxonomy, in terms of the choices made with regard to the search space (including choices that involve the covariates, if applicable) and the criterion function (including choices about the outcome variable(s) of interest). Subsequently, motivations for those choices will be presented.

Example 3.1. Hollander et al. (2017) examined the impact of the administration of canagliflozin (CANA) and phentermine (PHEN) on overweight and obese patients without type 2 diabetes. A clinical trial was conducted to compare four

15214036, 2023, 8, Downloaded from https://onlinelibrary.wiley.com/doi/10.1002/bimj.202200285 by <Shibboleth>-member@kuleuven.be, Wiley Online Library on [08/01/2024]. See the Terms and Conditions (https://onlinelibrary.wiley.com/terms-and-conditions) on Wiley Online Library for rules of use; OA articles are governed by the applicable Creative Commons License

TABLE 2 Choices with regard to search space and criterion function, estimated optimal treatment regimes ($\widehat{\text{OTR}}$), and associated estimated values of criterion function ($\widehat{E(Y_r^*)}$) for Example 3.1.

N°	Type of TR	Outcome	Criterion	$\widehat{\text{OTR}}$	$\widehat{E(Y_r^*)}$
1	One-size-fits-all	Y_1	Mean	$p \mapsto 3$	7.5
2	One-size-fits-all	Y_2	Mean	$p \mapsto 3$	0.667

Note. $\mathcal{A}$: 0 = Placebo, 1 = CANA, 2 = PHEN, 3 = CANA+PHEN;
Y_1 =Percentage change in body weight, Y_2 =weight loss $\geq$ 5%.

TABLE 3 Choices with regard to search space and criterion function, estimated optimal treatment regimes ($\widehat{\text{OTR}}$), and associated estimated values of criterion function ($\widehat{E(Y_r^*)}$) for Example 3.2.

N°	Type of TR	Function class	Outcome	Criterion	$\widehat{\text{OTR}}$	$\widehat{E(Y_r^*)}$
1	Precision	$I(\beta_0 + \beta_1 X_1 + \beta_2 \log(X_2 + 1) > 0)$	Y	Mean	$I(-60.0 + age + 7.98 \times \log(PR + 1) > 0)$	0.695
2	Precision	$1 - I(X_1 < \eta_0 \text{ and } X_2 < \eta_1)$	Y	Mean	$1 - I(age < 60 \text{ and } PR < 9)$	0.686

Note. $\mathcal{A}$: 0 = PF, 1 = PFT; X_1= age, X_2= PR; Y = disease-free survival to 3 years from baseline.
For identifiability purposes the weight of age was fixed at 1.

treatment alternatives: Placebo (0), CANA (1), PHEN (2), and CANA + PHEN (3). The primary endpoint was the percent change in body weight from baseline to week 26, and the secondary endpoint was whether or not the participant achieved a weight loss $\geq$ 5%. The results showed that the average percent change for the four treatment alternatives was 0.6%, 1.9%, 4.1%, and 7.5%, respectively, and that the proportion of patients with weight loss $\geq$ 5% was 17.5%, 17.9%, 41.7%, and 66.7%, respectively. Based on the results, the authors concluded that CANA + PHEN produces meaningful reductions in body weight and is the treatment of choice.

Regarding the *search space*, first, the chosen target set consisted of overweight and obese patients without type 2 diabetes. Second, as within that target set the four treatment alternatives were compared in terms of the resulting outcomes, this implies that a one-size-fits-all type of treatment regimes was chosen. Regarding the *criterion function*, first, two endpoints or outcome variables were chosen, with the secondary of those being obtained by thresholding the primary one—we will further deal with them as two variables involved in two different optimization exercises. Second, for both outcome variables, the distributional characteristic focused on in the criterion function was the mean. A schematic overview of all choices in question along with the associated results is given in Table 2.

The choices involved were motivated as follows: The choice of the target set was justified based on previous empirical evidence on the effectiveness of the administration of CANA to overweight and obese patients with and without type 2 diabetes. No explicit motivation was further given for the choice of a one-size-fits-all approach and mean-based criterion functions; both may probably be considered as default choices that go with an easy implementation and interpretation. Finally, the chosen type of preprocessing that led to the secondary endpoint is very typical in weight management studies, with the 5% threshold for weight loss from baseline generally being accepted as the threshold for "clinical meaningfulness" (Williamson et al., 2015).

Example 3.2. Zhang, Tsiatis, Laber, et al. (2012) reanalyzed data from a clinical trial comparing L-phenylalanine mustard plus 5-fluorouracil (PF, treatment alternative 0) with PF plus tamoxifen (PFT, treatment alternative 1) in patients with primary operable breast cancer and positive lymph nodes (Fisher et al., 1975, 1981, 1983). The raw endpoint was disease-free survival time, which was converted into 1 if a subject survived disease-free to 3 years from baseline, and 0 otherwise. Two kinds of treatment regimes were considered: (i) $I(\beta_0 + \beta_1 X_1 + \beta_2 \log(X_2 + 1) > 0)$ and (ii) $1 - I(X_1 < \eta_0 \text{ and } X_2 < \eta_1)$, where X_1 is age (in years) and X_2 is progesterone receptor level (PR, in fmol). This led to the following estimated OTRs: (i) $I(-60.0 + age + 7.98 \times \log(PR + 1) > 0)$ and (ii) $1 - I(age < 60 \text{ and } PR < 9)$. In spite of being different, both regimes imply that PF should be given to younger patients with low PR levels, and PFT to the other ones.

Regarding the *search space*, first, the chosen target set consisted of patients with primary operable breast cancer and positive lymph nodes. Second, a precision approach based on post hoc subgroups was chosen, with two prespecified classes of functions (viz., a linear and a specific tree-like one), and with within the linear class a log transformation being applied to one of the covariates. Regarding the *criterion function*, first, the authors decided to preprocess (viz., dichotomize) the raw endpoint. Second, they opted for a mean- (or probability-)based criterion function. A schematic overview of all choices in question along with the associated results is given in Table 3.

15214036, 2023, 8, Downloaded from https://onlinelibrary.wiley.com/doi/10.1002/bimj.202200285 by <Shibboleth>-member@kuleuven.be, Wiley Online Library on [08/01/2024]. See the Terms and Conditions (https://onlinelibrary.wiley.com/terms-and-conditions) on Wiley Online Library for rules of use; OA articles are governed by the applicable Creative Commons License

TABLE 4 Choices with regard to search space and criterion function, estimated optimal treatment regimes ($\widehat{\text{OTR}}$), and associated estimated values of criterion function ($\widehat{E(Y_r^*)}$, $\widehat{Q_{.25}(Y_r^*)}$) for Example 3.3.

N°	Type of TR	Function class	Outcome	Criterion	$\widehat{\text{OTR}}$	$\widehat{E(Y_r^*)}$	$\widehat{Q_{.25}(Y_r^*)}$
1	Precision	$I(\eta_0 + \eta_1 X_1 + \eta_2 X_2 > 0)$	Y	Mean	$I(0.53 - 0.80X_1 - 0.29X_2) > 0)$	403.9	219
2	Precision	$I(\eta_0 + \eta_1 X_1 + \eta_2 X_2 > 0)$	Y	Quantile .25	$I(0.21 - 0.96X_1 + 0.19X_2) > 0)$	346.5	263

Note. $\mathcal{A}$: 0 = ddI, 1 = ddI+AZT; X_1 = weight at baseline, X_2 = CD4 T cell count at baseline; Y = CD4 T cell count at week 96.

The choices involved could be motivated as follows: The choice of the target set of patients was based on the hypothesis that the chemotherapeutic agents under study might work for it (Fisher et al., 1975, 1981, 1983). The choice of a precision type of TR was further motivated by Zhang, Tsiatis, Laber, et al. (2012) by referring to empirical evidence of heterogeneity in response to PFT therapy depending on age and progesterone receptor level (Fisher et al., 1983; Gail & Simon, 1985). The two prespecified classes of functions were further chosen because of their simplicity and ease of interpretation (with, for the first class, the covariate PR being preprocessed by means of a log transformation because of its extreme skewness); moreover, the second class was also chosen for reasons of comparison with previous results obtained with the same data (Fisher et al., 1983; Gail & Simon, 1985). Finally, as the raw endpoint was subject to censoring, some preprocessing of it was desirable (as discussed in Section 2.4). Given the choice to do so by means of a dichotomization, the subsequent choice of a probability-based criterion function was pretty obvious.

Example 3.3. Wang et al. (2018) reanalyzed data from a clinical trial (Hammer et al., 1996) comparing the administration of didanosine (ddI, treatment alternative 0) with that of didanosine + azidothymidine (ddI+AZT, treatment alternative 1) to HIV-infected patients (who already took AZT before entering the trial). The response variable was CD4 T cell count at week 96. The authors considered treatment regimes of the type $I(\eta_0 + \eta_1 X_1 + \eta_2 X_2 > 0)$, with X_1 denoting the patient's weight (kg) at baseline and X_2 the patient's baseline CD4 T cell count (with both covariates being standardized to the interval $[0, 1]$). Two criteria, mean and quantile, were considered, resulting in two estimated OTRs; these appeared to be rather different, especially in terms of the weight of X_1.

Regarding the *search space*, first, the chosen target set essentially consists of the HIV patients. Second, a precision approach with post hoc subgroups was chosen, along with a prespecification of a linear class of functions involving two covariates. Regarding the *criterion function*, two different distributional characteristics were considered as a basis of two criterion functions to estimate an OTR. A schematic overview of all choices in question, along with the associated results, is given in Table 4.

Regarding the basis of the above choices, the choice of a precision type of TR was motivated, in part, by previous empirical evidence on sizable individual differences in the pharmacokinetics of AZT, which were related to, among other things, body weight (Hammer et al., 1996); otherwise, this also justified the inclusion of baseline weight, in addition to baseline CD4 cell count, as covariates. A linear class of functions was selected because of its simplicity and interpretability. Finally, next to the mean criterion as a benchmark, a quantile criterion was chosen as the distribution of the outcome variable (CD4 count at week 96) tended to be asymmetric and skewed to the right (with outliers in the right tail possibly obfuscating the mean criterion), and because of the desire to improve the treatment effect for weaker patients.

4 | GENERAL DISCUSSION

Motivated by the fact that the concept of optimality of decision rules in general and treatment regimes in particular has received so far relatively little attention and discussion, we presented reflections on its conceptual basis. We did so starting from the basics of any optimization problem. Subsequently, we considered at length links with both formal points of interest (such as, e.g., respect for the nature of the variables involved, and replicability of the results) and substantive concerns (such as, e.g., sound theoretical justifications and interpretations, and capturing the core aspects practitioners care about). This conceptual basis may, on the one hand, help sharpen statistical thinking about optimality of decision rules for treatment assignment (also linking up with the first statistical virtue advanced by Gelman & Hennig, 2017, with regard to clear and unambiguous definitions of concepts). On the other hand, it may also facilitate the dialogue between the statistician or statistical consultant and the applied researcher or practitioner in search of an OTR. Indeed, regarding the latter, it may provide a frame of reference, within which the initial concern of the applied researcher can be adequately captured, and subsequently sharpened and broadened. Moreover, the conceptual basis may also help in highlighting the distinctive

15214036, 2023, 8, Downloaded from https://onlinelibrary.wiley.com/doi/10.1002/bimj.202200285 by <Shibboleth>-member@kuleuven.be, Wiley Online Library on [08/01/2024]. See the Terms and Conditions (https://onlinelibrary.wiley.com/terms-and-conditions) on Wiley Online Library for rules of use; OA articles are governed by the applicable Creative Commons License

contributions of each of the partners to this dialogue. Specifically, the statistician may rely on a deeper understanding of formal aspects and properties (such as, e.g., nestedness of search spaces and the precise meaning of distributional characteristics), and may contribute to a faithful translation of clinical concerns into mathematical/statistical terms. By contrast, the clinician may contribute key substantive (e.g., etiological, pharmacokinetic) theories and hypotheses, expert knowledge on relevant results from previous empirical studies, and vital pragmatic information such as the cost of collecting particular types of covariate information (e.g., information on biomarkers); likewise, the clinician is uniquely qualified to put forward substantive concerns (e.g., a concern about the fraction of patients suffering from adverse treatment effects).

As discussed at length in this paper, at many different points in the process of specifying the type of optimality of treatment regimes that bears on a particular application, several choices need to be made. All of these choices require a careful reflection about relevant formal and substantive concerns. This especially also holds because the choices in question are not without consequences, as illustrated by both synthetic and empirical examples (viz., Examples 2.1 and 3.3). That being said, it may often be meaningful to consider multiple choice alternatives at several points. There are at least three reasons for this:

1. Even in case substantive concerns dictate a particular type of optimality, it may not be harmful to consider alternative specifications, which may lead to useful additional information (also in line with the fifth statistical virtue suggested by Gelman & Hennig, 2017, "awareness of multiple perspectives"). As an example, even if one wishes to go for a one-size-fits-all regime, exploring a precision approach may not be harmful, in that it nonetheless may yield a one-size-fits-all regime (because of the usual nestedness of one-size-fits-all and precision search spaces); moreover, a precision approach may yield important additional information, such as the identification of subgroups of patients who do not benefit from the optimal one-size-fits-all regime (or, stated in other words, whereas the one-size-fits-all regime may constitute the "rule," there may be "exceptions to the rule"). As a second example, even if one wishes to go for a purely confirmatory, hypothesis-testing approach, some concomitant explorations may allow the data to unveil unexpected and useful additional information.
2. In case trying out several choice options (e.g., different families of functions in a precision approach) leads to converging results (as in Example 3.2), this may contribute to more robust and replicable conclusions. In general, such a method of working has obvious links with the principles underlying sensitivity (Saltelli et al., 2008) and multiverse (Steegen et al., 2016) analyses.
3. Especially useful information could be derived from searching optimal regimes that are obtained by specifying different search spaces while keeping the target set of patients and the criterion function constant, and from subsequently comparing the resulting optimal regimes in terms of their value on the criterion function, $C(Y^*_{r^{opt}})$. In particular, this may allow to answer questions such as: Which gain could be obtained by replacing the current gold standard regime in clinical practice by a newly estimated OTR, or by replacing a one-size-fits-all OTR by a precision OTR? Is that gain sufficiently large to warrant the replacement in question—with the option to factor in at this point a threshold for an improvement to be considered clinically significant (Friedman et al., 2015; Gad, 2009; Spiegelhalter et al., 1994; Tang & Tu, 2012)? Still another group of questions is of relevance if, within a precision approach, different families of functions are considered; one may then wonder to what extent possible gains for more complex families, in terms of higher values of $C(Y^*_{r^{opt}})$, come at the expense of a loss of interpretability/insightfulness, in view of striking an optimal balance between both. As a third example of a group of questions, one may wish to examine whether the inclusion in a precision regime of a number of covariates for which the collection of information is difficult or expensive pays off in terms of a clinically significant increase in the value of the criterion function.

Of course this paper has some limitations. First, we focused on decision rules with a single decision point only. This, however does not seem to be a real restriction in that a generalization to a context with multiple decision points (i.e., the context of dynamic treatment regimes Tsiatis et al., 2020) is pretty straightforward.

Second, we exclusively focused on settings that involve a single (scalar-valued) outcome variable only. Applied contexts, however, not seldom include multiple outcome variables (which may pertain to both beneficial and adverse treatment effects). So far, a number of approaches have been proposed to deal with multiple outcome variables: combining them into one single outcome variable (Bakal et al., 2015; Myles et al., 2014; Tong et al., 2012), considering a single variable as the main outcome and the others as additional ones (Imai & Strauss, 2011; Laber et al., 2018), and considering each of the outcomes separately and subsequently chart in which areas of the covariate space regimes yield conflicting assignments (Laber, Lizotte, & Ferguson, 2014). Each of these approaches goes with distinctive questions to be dealt with (such as, e.g., how to combine outcome variables that have been typically measured on different scales). Importantly, however,

15214036, 2023, 8, Downloaded from https://onlinelibrary.wiley.com/doi/10.1002/bimj.202200285 by <Shibboleth>-member@kuleuven.be, Wiley Online Library on [08/01/2024]. See the Terms and Conditions (https://onlinelibrary.wiley.com/terms-and-conditions) on Wiley Online Library for rules of use; OA articles are governed by the applicable Creative Commons License

from the viewpoint of this paper, the conceptual basis proposed in it can be largely extended to each of the approaches outlined above.

Finally, an important aspect of the identification of an OTR is the actual optimization or OTR estimation, which we left out of consideration in this paper. For this purpose, we refer the interested reader once again to the books and review articles mentioned in Section 1. It should be noted that, at present, estimation algorithms are not yet available for all possible optimization problems that could be conceived with the conceptual basis outlined above. A more fundamental problem is that the criterion function that is to be optimized in the search of an OTR is, in general, typically defined on the level of the population (as $C(Y^*_{r^{opt}})$ in (1)). In practice, however, the actual optimization or OTR estimation has to take place on the level of the sample data. This implies that, in the optimization, an estimator of the population criterion function will be needed. This issue will go with questions on the quality of the estimator that may be picked for this purpose. Otherwise, further advancing this issue, also making correct inferences about the quality on the level of the population of an estimated OTR (Sies & Van Mechelen, 2019) (as well as about the difference in such quality between different estimated OTR's) is a critical inferential issue, which, in part, is an open problem that may be considered a challenge for further statistical research.

5 | CONCLUDING REMARK

In many areas, applied researchers as well practitioners have to choose between different solutions for a problem at hand. This immediately calls for optimal decision rules to settle the choices involved. A derivation of valid such rules requires a careful specification of the type of optimality one cares about. In this paper, we presented reflections on the concept of optimality that may provide a useful hold in this specification. That being said, we wish that the presented reflections and the associated discussion may also contribute to an improved accessibility of the domain of optimal decision rules, which will hopefully further lead to more and convincing substantive applications.

ACKNOWLEDGMENTS

The research was supported in part by the Research Foundation - Flanders (G080219N) and by the Research Fund of KU Leuven (C14/19/054).

CONFLICT OF INTEREST STATEMENT

The authors have declared no conflict of interest.

DATA AVAILABILITY STATEMENT

The data that support the findings of this study are available in the Supporting Information of this paper.

OPEN RESEARCH BADGES

This article has earned an Open Data badge for making publicly available the digitally-shareable data necessary to reproduce the reported results. The data is available in the Supporting Information section.

This article has earned an open data badge "**Reproducible Research**" for making publicly available the code necessary to reproduce the reported results. The results reported in this article could fully be reproduced.

ORCID

Trung Dung Tran https://orcid.org/0000-0001-9967-1012

REFERENCES

Ascarza, E., Neslin, S. A., Netzer, O., Anderson, Z., Fader, P. S., Gupta, S., Hardie, B. G. S., Lemmens, A., Libai, B., Neal, D., Provost, F., & Schrift, R. (2018). In pursuit of enhanced customer retention management: Review, key issues, and future directions. *Customer Needs and Solutions*, *5*(1), 65–81.

Bakal, J. A., Westerhout, C. M., & Armstrong, P. W. (2015). Impact of weighted composite compared to traditional composite endpoints for the design of randomized controlled trials. *Statistical Methods in Medical Research*, *24*(6), 980–988.

Bose, G., & Freedman, M. S. (2020). Precision medicine in the multiple sclerosis clinic: Selecting the right patient for the right treatment. *Multiple Sclerosis Journal*, *26*(5), 540–547.

15214036, 2023, 8, Downloaded from https://onlinelibrary.wiley.com/doi/10.1002/bimj.202200285 by <Shibboleth>-member@kuleuven.be, Wiley Online Library on [08/01/2024]. See the Terms and Conditions (https://onlinelibrary.wiley.com/terms-and-conditions) on Wiley Online Library for rules of use; OA articles are governed by the applicable Creative Commons License

Brinkley, J., Tsiatis, A., & Anstrom, K. J. (2010). A generalized estimator of the attributable benefit of an optimal treatment regime. *Biometrics*, *66*(2), 512–522.

Brookes, S. T., Whitely, E., Egger, M., Smith, G. D., Mulheran, P. A., & Peters, T. J. (2004). Subgroup analyses in randomized trials: Risks of subgroup-specific analyses; power and sample size for the interaction test. *Journal of Clinical Epidemiology*, *57*(3), 229–236.

Caminada, K., Goudswaard, K., & Wang, C.-K. (2012). *Disentangling income inequality and the redistributive effect of taxes and transfers in 20 LIS countries over time*. LIS Working Paper, 581.

Chakraborty, B., & Moodie, E. (2013). *Statistical methods for dynamic treatment regimes: Reinforcement learning, causal inference, and personalized medicine*. Springer.

Chakraborty, B., & Murphy, S. A. (2014). Dynamic treatment regimes. *Annual Review of Statistics and Its Application*, *1*(1), 447–464.

Demonty, I., Ras, R. T., van der Knaap, H. C. M., Duchateau, G. S. M. J. E., Meijer, L., Zock, P. L., Geleijnse, J. M., & Trautwein, E. A. (2009). Continuous dose-response relationship of the LDL-cholesterol-lowering effect of phytosterol intake. *The Journal of Nutrition*, *139*(2), 271–284.

Doestzada, M., Vila, A. V., Zhernakova, A., Koonen, D. P. Y., Weersma, R. K., Touw, D. J., Kuipers, F., Wijmenga, C., & Fu, J. (2018). Pharmacomicrobiomics: A novel route towards personalized medicine? *Protein & Cell*, *9*(5), 432–445.

Dusseldorp, E., & Van Mechelen, I. (2014). Qualitative interaction trees: A tool to identify qualitative treatment-subgroup interactions. *Statistics in Medicine*, *33*(2), 219–237.

Feinstein, A. R. (1998). The problem of cogent subgroups: A clinicostatistical tragedy. *Journal of Clinical Epidemiology*, *51*(4), 297–299.

Fisher, B., Carbone, P., Economou, S. G., Frelick, R., Glass, A., Lerner, H., Redmond, C., Zelen, M., Band, P., Katrych, D. L., Wolmark, N., & Fisher, E. R. (1975). L-phenylalanine mustard (L-PAM) in the management of primary breast cancer. *New England Journal of Medicine*, *292*(3), 117–122.

Fisher, B., Redmond, C., Brown, A., Wickerham, D. L., Wolmark, N., Allegra, J., Escher, G., Lippman, M., Savlov, E., & Wittliff, J. (1983). Influence of tumor estrogen and progesterone receptor levels on the response to tamoxifen and chemotherapy in primary breast cancer. *Journal of Clinical Oncology*, *1*(4), 227–241.

Fisher, B., Redmond, C., Brown, A., Wolmark, N., Wittliff, J., Fisher, E. R., Plotkin, D., Bowman, D., Sachs, S., Wolter, J., Frelick, R., Desser, R., LiCalzi, N., Geggie, P., Campbell, T., Elias, E. G., Prager, D., Koontz, P., Volk, H., & Shibata, H. (1981). Treatment of primary breast cancer with chemotherapy and tamoxifen. *New England Journal of Medicine*, *305*(1), 1–6.

Flashner-Abramson, E., Vasudevan, S., Adejumobi, I. A., Sonnenblick, A., & Kravchenko-Balasha, N. (2019). Decoding cancer heterogeneity: Studying patient-specific signaling signatures towards personalized cancer therapy. *Theranostics*, *9*(18), 5149–5165.

Foster, J. C., Taylor, J. M. G., & Ruberg, S. J. (2011). Subgroup identification from randomized clinical trial data. *Statistics in Medicine*, *30*(24), 2867–2880.

Friedman, L., Furberg, C., DeMets, D., Reboussin, D., & Granger, C. (2015). *Fundamentals of clinical trials*. Springer.

Gad, S. (2009). *Clinical trials handbook*. Wiley.

Gail, M., & Simon, R. (1985). Testing for qualitative interactions between treatment effects and patient subsets. *Biometrics*, *41*(2), 361–372.

Gelman, A., & Hennig, C. (2017). Beyond subjective and objective in statistics. *Journal of the Royal Statistical Society Series A: Statistics in Society*, *180*(4), 967–1033.

Gerstung, M., Papaemmanuil, E., Martincorena, I., Bullinger, L., Gaidzik, V. I., Paschka, P., Heuser, M., Thol, F., Bolli, N., Ganly, P., Ganser, A., McDermott, U., Döhner, K., Schlenk, R. F., Döhner, H., & Campbell, P. J. (2017). Precision oncology for acute myeloid leukemia using a knowledge bank approach. *Nature Genetics*, *49*(3), 332–340.

Goldberg, Y., & Kosorok, M. R. (2012). Q-learning with censored data. *The Annals of Statistics*, *40*(1), 529–560.

Hahn, P. R., Murray, J. S., & Carvalho, C. M. (2020). Bayesian regression tree models for causal inference: Regularization, confounding, and heterogeneous effects (with discussion). *Bayesian Analysis*, *15*(3), 965–1056.

Hammer, S. M., Katzenstein, D. A., Hughes, M. D., Gundacker, H., Schooley, R. T., Haubrich, R. H., Henry, W. K., Lederman, M. M., Phair, J. P., Niu, M., Hirsch, M. S., & Merigan, T. C. (1996). A trial comparing nucleoside monotherapy with combination therapy in HIV-infected adults with cd4 cell counts from 200 to 500 per cubic millimeter. *New England Journal of Medicine*, *335*(15), 1081–1090.

Hanushek, E. A., & Wäßmann, L. (2006). Does educational tracking affect performance and inequality? Differences-in-differences evidence across countries. *The Economic Journal*, *116*(510), C63–C76.

Hollander, P., Bays, H. E., Rosenstock, J., Frustaci, M. E., Fung, A., Vercruysse, F., & Erondu, N. (2017). Coadministration of canagliflozin and phentermine for weight management in overweight and obese individuals without diabetes: A randomized clinical trial. *Diabetes Care*, *40*(5), 632–639.

Imai, K., & Ratkovic, M. (2013). Estimating treatment effect heterogeneity in randomized program evaluation. *Annals of Applied Statistics*, *7*(1), 443–470.

Imai, K., & Strauss, A. (2011). Estimation of heterogeneous treatment effects from randomized experiments, with application to the optimal planning of the get-out-the-vote campaign. *Political Analysis*, *19*(1), 1–19.

Imbens, G., & Rubin, D. (2015). *Causal inference for statistics, social, and biomedical sciences: An introduction*. Cambridge University Press.

Jiang, R., Lu, W., Song, R., & Davidian, M. (2017). On estimation of optimal treatment regimes for maximizing t-year survival probability. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, *79*(4), 1165–1185.

Kiely, P. D. W. (2015). Biologic efficacy optimization-a step towards personalized medicine. *Rheumatology*, *55*(5), 780–788.

Kosorok, M., & Moodie, E. (2015). *Adaptive treatment strategies in practice: Planning trials and analyzing data for personalized medicine*. Society for Industrial and Applied Mathematics.

15214036, 2023, 8, Downloaded from https://onlinelibrary.wiley.com/doi/10.1002/bimj.202200285 by <Shibboleth>-member@kuleuven.be, Wiley Online Library on [08/01/2024]. See the Terms and Conditions (https://onlinelibrary.wiley.com/terms-and-conditions) on Wiley Online Library for rules of use; OA articles are governed by the applicable Creative Commons License

Kurian, A. W., Lichtensztajn, D. Y., Keegan, T. H. M., Nelson, D. O., Clarke, C. A., & Gomez, S. L. (2014). Use of and mortality after bilateral mastectomy compared with other surgical treatments for breast cancer in California, 1998-2011. *JAMA*, *312*(9), 902–914.

Laber, E., Lizotte, D. J., & Ferguson, B. (2014). Set-valued dynamic treatment regimes for competing outcomes. *Biometrics*, *70*(1), 53–61.

Laber, E., Lizotte, D. J., Qian, M., Pelham, W. E., & Murphy, S. A. (2014). Dynamic treatment regimes: Technical challenges and applications. *Electronic Journal of Statistics*, *8*(1), 1225–1272.

Laber, E., Wu, F., Munera, C., Lipkovich, I., Colucci, S., & Ripa, S. (2018). Identifying optimal dosage regimes under safety constraints: An application to long term opioid treatment of chronic pain. *Statistics in Medicine*, *37*(9), 1407–1418.

Laber, E., & Zhao, Y. Q. (2015). Tree-based methods for individualized treatment regimes. *Biometrika*, *102*(3), 501–514.

Lee, S., Lei, M., & Brody, G. (2015). Confidence intervals for distinguishing ordinal and disordinal interactions in multiple regression. *Psychological Methods*, *20*(2), 245–258.

Linn, K. A., Laber, E., & Stefanski, L. A. (2017). Interactive Q-learning for quantiles. *Journal of the American Statistical Association*, *112*(518), 638–649.

Lipkovich, I., Dmitrienko, A., & Agostino, R. (2017). Tutorial in biostatistics: Data-driven subgroup identification and analysis in clinical trials. *Statistics in Medicine*, *36*(1), 136–196.

Myles, P. S., Leslie, K., Chan, M. T. V., Forbes, A., Peyton, P. J., Paech, M. J., Beattie, W. S., Sessler, D. I., Devereaux, P. J., Silbert, B., Schricker, T., & Wallace, S. (2014). The safety of addition of nitrous oxide to general anaesthesia in at-risk patients having major non-cardiac surgery (enigma-ii): A randomised, single-blind trial. *The Lancet*, *384*(9952), 1446–1454.

National Research Council. (2011). *Toward precision medicine: Building a knowledge network for biomedical research and a new taxonomy of disease*. The National Academies Press.

Neyman, J. (1923). On the application of probability theory to agricultural experiments: Essay on principles (in Polish). *Roczniki Nauk Rolniczych*, *10*, 1–51.

Nicholson, R., Griffiths, K., Blamey, R., Bishop, H., & Haybittle, J. (1981). Estrogen receptors and breast cancer. *Environmental Health Perspectives*, *38*, 143–146.

Qian, M., & Murphy, S. A. (2011). Performance guarantees for individualized treatment rules. *Annals of Statistics*, *39*(2), 1180–1210.

Rothwell, P. M. (2005). Subgroup analysis in randomised controlled trials: Importance, indications, and interpretation. *The Lancet*, *365*(9454), 176–186.

Rubin, D. B. (1974). Estimating causal effects of treatments in randomized and nonrandomized studies. *Journal of Educational Psychology*, *66*(5), 688–701.

Rubin, D. B. (1978). Bayesian inference for causal effects: The role of randomization. *The Annals of Statistics*, *6*(1), 34–58.

Rubin, D. B. (2005). Causal inference using potential outcomes. *Journal of the American Statistical Association*, *100*(469), 322–331.

Saltelli, A., Ratto, M., Andres, T., Campolongo, F., Cariboni, J., Gatelli, D., Saisana, M., & Tarantola, S. (2008). *Global sensitivity analysis: The primer*. Wiley.

Saraç, L., & Ok, A. (2010). The effects of different instructional methods on students' acquisition and retention of cardiopulmonary resuscitation skills. *Resuscitation*, *81*(5), 555–561.

Schulte, P. J., Tsiatis, A., Laber, E., & Davidian, M. (2014). Q-and A-learning methods for estimating optimal dynamic treatment regimes. *Statistical Science*, *29*(4), 640–661.

Sies, A., & Van Mechelen, I. (2017). Comparing four methods for estimating tree-based treatment regimes. *The International Journal of Biostatistics*, *13*(1), 20160068.

Sies, A., & Van Mechelen, I. (2019). Estimating the quality of optimal treatment regimes. *Statistics in Medicine*, *38*(25), 4925–4938.

Spiegelhalter, D. J., Freedman, L. S., & Parmar, M. K. B. (1994). Bayesian approaches to randomized trials. *Journal of the Royal Statistical Society: Series A (Statistics in Society)*, *157*(3), 357–387.

Steegen, S., Tuerlinckx, F., Gelman, A., & Vanpaemel, W. (2016). Increasing transparency through a multiverse analysis. *Perspectives on Psychological Science*, *11*(5), 702–712.

Tang, W., & Tu, X. (2012). *Modern clinical trial analysis*. Springer.

Tong, B. C., Huber, J. C., Ascheim, D. D., Puskas, J. D., Ferguson, T. B., Blackstone, E. H., & Smith, P. K. (2012). Weighting composite endpoints in clinical trials: Essential evidence for the heart team. *The Annals of Thoracic Surgery*, *94*(6), 1908–1913.

Trivedi, M. H. (2016). Right patient, right treatment, right time: Biosignatures and precision medicine in depression. *World Psychiatry*, *15*(3), 237–238.

Tsiatis, A., Davidian, M., Holloway, S., & Laber, E. (2020). *Dynamic treatment regimes: Statistical methods for precision medicine*. Chapman and Hall/CRC.

Wager, S., & Athey, S. (2018). Estimation and inference of heterogeneous treatment effects using random forests. *Journal of the American Statistical Association*, *113*(523), 1228–1242.

Wang, L., Zhou, Y., Song, R., & Sherwood, B. (2018). Quantile-optimal treatment regimes. *Journal of the American Statistical Association*, *113*(523), 1243–1254.

Whitcomb, D. C. (2019). Primer on precision medicine for complex chronic disorders. *Clinical and Translational Gastroenterology*, *10*(7), e00067–e00067.

Williamson, D. A., Bray, G. A., & Ryan, D. H. (2015). Is 5% weight loss a satisfactory criterion to define clinically significant weight loss? *Obesity*, *23*(12), 2319–2320.

15214036, 2023, 8, Downloaded from https://onlinelibrary.wiley.com/doi/10.1002/bimj.202200285 by <Shibboleth>-member@kuleuven.be, Wiley Online Library on [08/01/2024]. See the Terms and Conditions (https://onlinelibrary.wiley.com/terms-and-conditions) on Wiley Online Library for rules of use; OA articles are governed by the applicable Creative Commons License

Wright, J. M., Musini, V. M., & Gill, R. (2018). First-line drugs for hypertension. *The Cochrane Database of Systematic Reviews*, *4*(4), CD001841–CD001841.

Xu, Y., Yu, M., Zhao, Y., Li, Q., Wang, S., & Shao, J. (2015). Regularized outcome weighted subgroup identification for differential treatment effects. *Biometrics*, *71*(3), 645–653.

Zajonc, T. (2012). Bayesian inference for dynamic treatment regimes: Mobility, equity, and efficiency in student tracking. *Journal of the American Statistical Association*, *107*(497), 80–92.

Zhang, B., Tsiatis, A., Davidian, M., Zhang, M., & Laber, E. (2012). Estimating optimal treatment regimes from a classification perspective. *Stat*, *1*(1), 103–114.

Zhang, B., Tsiatis, A., Laber, E., & Davidian, M. (2012). A robust method for estimating optimal treatment regimes. *Biometrics*, *68*(4), 1010–1018.

Zhang, B., Tsiatis, A., Laber, E., & Davidian, M. (2013). Robust estimation of optimal dynamic treatment regimes for sequential treatment decisions. *Biometrika*, *100*(3), 681–694.

Zhang, Y., Laber, E., Tsiatis, A., & Davidian, M. (2015). Using decision lists to construct interpretable and parsimonious treatment regimes. *Biometrics*, *71*(4), 895–904.

Zhao, Y., Kosorok, M. R., & Zeng, D. (2009). Reinforcement learning design for cancer clinical trials. *Statistics in Medicine*, *28*(26), 3294–3315.

Zhao, Y., Zeng, D., Rush, A. J., & Kosorok, M. R. (2012). Estimating individualized treatment rules using outcome weighted learning. *Journal of the American Statistical Association*, *107*(499), 1106–1118.

How to cite this article: Tran, T. D., Abad, A. A., Verbeke, G., Molenberghs, G., & Van Mechelen, I. (2023). Reflections on the concept of optimality of single decision point treatment regimes. *Biometrical Journal*, *65,* 2200285. https://doi.org/10.1002/bimj.202200285

15214036, 2023, 8, Downloaded from https://onlinelibrary.wiley.com/doi/10.1002/bimj.202200285 by <Shibboleth>-member@kuleuven.be, Wiley Online Library on [08/01/2024]. See the Terms and Conditions (https://onlinelibrary.wiley.com/terms-and-conditions) on Wiley Online Library for rules of use; OA articles are governed by the applicable Creative Commons License