Original Research Article

MDM Medical Decision Making

Medical Decision Making
2022, Vol. 42(5) 649–660

DOI: 10.1177/0272989X211072858
journals.sagepub.com/home/mdm
SAGE

Prioritizing Research in an Era of Personalized Medicine: The Potential Value of Unexplained Heterogeneity

Anna Heath, and Petros Pechlivanoglou

Background. Clinical care is moving from a "one size fits all" approach to a setting in which treatment decisions are based on individual treatment response, needs, preferences, and risk. Research into personalized treatment strategies aims to discover currently unknown markers that identify individuals who would benefit from treatments that are nonoptimal at the population level. Before investing in research to identify these markers, it is important to assess whether such research has the potential to generate value. Thus, this article aims to develop a framework to prioritize research into the development of new personalized treatment strategies by creating a set of measures that assess the value of personalizing care based on a set of unknown patient characteristics. **Methods.** Generalizing ideas from the value of heterogeneity framework, we demonstrate 3 measures that assess the value of developing personalized treatment strategies. The first measure identifies the potential value of personalizing medicine within a given disease area. The next 2 measures highlight specific research priorities and subgroup structures that would lead to improved patient outcomes from the personalization of treatment decisions. **Results.** We graphically present the 3 measures to perform sensitivity analyses around the key drivers of value, in particular, the correlation between the individual treatment benefits across the available treatment options. We illustrate these 3 measures using a previously published decision model and discuss how they can direct research in personalized medicine. **Conclusion.** We discuss 3 measures that form the basis of a novel framework to prioritize research into novel personalized treatment strategies. Our novel framework ensures that research targets personalized treatment strategies that have high potential to improve patient outcomes and health system efficiency.

Highlights

- It is important to undertake research prioritization before conducting any research that aims to discover novel methods (e.g., biomarkers) for personalizing treatment.
- The value of unexplained heterogeneity can highlight disease areas in which personalizing treatment can be valuable and determine key priorities within that area.
- These priorities can be determined under assumptions of the magnitude of the individual-level treatment effect, which we explore in sensitivity analyses.

Keywords
precision medicine, personalized medicine, research prioritization, study design, simulation modeling, value of heterogeneity, value of information

Date received: May 13, 2021; accepted: December 20, 2021

Corresponding Author
Anna Heath, Child Health Evaluative Sciences, The Hospital for Sick Children, 686 Bay Street, Floor 11, L4 East, Toronto, ON M5G 0A4, Canada; (anna.heath@sickkids.ca).

Decision modeling in health generally aims to support population-level decision making by synthesizing the available evidence to identify the treatment strategy that offers the maximum benefit at the population level, among a set of potential options.[1] Although these "population-level" decisions are usually based on a small number of patient-level features (e.g., their diagnosis), greater clinical and/or economic value could be generated by personalizing treatment decisions.[2–4] This is because each individual would be offered the treatment that is expected to maximize the value associated with their treatment,[2] implying that an individual would be switched from the population-level optimal treatment strategy only if greater benefit could be derived from an alternative treatment option. However, to make accurate individual-level treatment decisions and realize this additional value, significant investment must be made in research to develop these novel personalized treatment strategies.

With limited budgets available, it is important to prioritize the funding of research that has the greatest potential to efficiently improve health outcomes, thereby increasing value in the health care system.[5] Thus, research into novel personalized treatment strategies should be directed toward disease areas in which making treatment decisions at the individual level has the potential to generate value. Once a suitable disease area has been identified (e.g., breast cancer), research can be further prioritized by focusing on the development of individualized treatment strategies that have the potential to generate substantial benefit. This article aims to develop 3 measures that can indicate, alongside subject-specific expertise, key research priorities that would allow the development of valuable personalized treatment strategies.

Value-of-information (VoI) methods have long been suggested as a method for research prioritization,[6–11] as they assess the impact of statistical uncertainty on decision making and prioritize research that efficiently reduces this uncertainty.[12] VoI methods require decision models that estimate the population-average benefit of each potential treatment strategy for a given disease. Typically, this population-average benefit is defined using the net monetary or net health benefit function,[13] where the net benefit values are calculated conditional on a set of decision model parameters. Uncertainty in the net benefit, and in the decision making, is induced by statistical uncertainty in the model parameters and is usually estimated by simulation. VoI then determines the value of collecting additional data to inform these parameters.[10]

VoI analyses assume that decision makers are searching for a single optimal treatment strategy to implement across the whole population of interest, and decision uncertainty arises from imperfect knowledge of model parameters.[14,15] However, VoI concepts have been extended to calculate the value of individualizing care by exploiting heterogeneity in individual patient response.[2,4,16] First, the expected value of individualized care (EVIC) calculates the value of personalizing care based on patient preferences.[2] To achieve this, patient preferences are valued (e.g., using quality-of-life [QoL] weights) and assumed to vary across the population, making the assumption that patients are homogeneous except for these differences in preference. EVIC can then be extended and combined with previously published concepts,[4] to define the value of heterogeneity (VoH) framework. VoH is a unified theory that quantifies the expected value to be gained from making stratified treatment decisions based on patient characteristics and the value of resolving parameter uncertainty within these subgroups.[16]

To calculate EVIC and VoH, we must estimate the net monetary benefit for each treatment option for each individual to determine the optimal treatment at the individual level. Once the individual optimal treatment is found, the value of personalizing treatment can be computed by comparing to the value of a single treatment for the population.[16] The EVIC calculations assume that patient preferences for a given health state are unaffected by treatment and an individual's net benefit for each treatment option can be computed using the value they assign to different health states. VoH calculations assume that treatment decisions are stratified based on known patient characteristics (e.g., gender, health status, or body mass index).[16] Individual-level data on treatment response and baseline characteristics can then be used to calculate the net benefits for each treatment, conditional on these known characteristics (e.g., using regression).[16]

However, research into individualized treatment strategies is often concerned with identifying new patient

Child Health Evaluative Sciences, The Hospital for Sick Children, Toronto, ON, Canada (AH, PP); Division of Biostatistics, Dalla Lana School of Public Health, University of Toronto, Toronto, ON, Canada (AH); Department of Statistical Science, University College London, London, UK (AH); and Institute of Health Policy, Management and Evaluation, University of Toronto, Toronto, ON, Canada (PP). The authors declared no potential conflicts of interest with respect to the research, authorship, and/or publication of this article. The authors disclosed receipt of the following financial support for the research, authorship, and/or publication of this article: Anna Heath was funded by the Canadian Institutes of Health Research through the PERC SPOR iPCT grant (funding No. MYG-151207; 2017–2020), a Canada Research Chair in Statistical Trial Design and the Natural Sciences and Engineering Research Council of Canada (award No. RGPIN-2021-03366). The funding agreement ensured the authors' independence in designing the study, interpreting the data, writing, and publishing the report.

subgroups that are defined by some, currently unknown, characteristic (e.g., a novel genetic or biologic marker or a clinical algorith).[17–19] In these instances, the characteristic that allows for individual treatment decisions is currently unknown. Thus, the individual net benefit across each treatment cannot be computed for a given individual, and the individual-level optimal treatment is unknown. This means that the current VoH measures cannot prioritize research studies that aim to identify these novel markers.

To address this, we generalize the VoH and EVIC measures to allow for the prioritization of research into novel personalized medicine strategies. We discuss 3 measures to generalize VoH and include a consideration of the value of unexplained heterogeneity (i.e., heterogeneity in the individual responses that cannot be explained by currently known characteristics). We demonstrate these measures using an individual-level decision model that we have adapted from a previously published model.[10] This article begins by introducing this model and its key assumptions. We then define each of the 3 measures in turn and discuss how they can help prioritize research. The first measure determines whether there is any potential to generate value by personalizing medicine in the disease area under investigation. Following this, the next 2 measures aim to highlight research areas where value is likely to be generated through the development of specific personalized treatment strategies. We discuss how these measures should be combined with subject matter expertise to undertake the research prioritization. We conclude with a discussion on the limitations of these proposed measures and suggestions for potential extensions.

Decision Making with Individual-Level Decision Models

The research prioritization framework developed in this article requires estimates of the individual-level net monetary/health benefit, a summary measure of the value of a treatment, measured in monetary or health units, respectively.[13] We assume that, in a given context, individuals are choosing between multiple treatment options. The heterogeneity in the individual-level net monetary or health benefit under each of the treatment options is estimated using an individual-level decision model that combines key outcomes into a single measure of net monetary or health benefit. Thus, the VoH is based on a model that incorporates individual-level variation into the net benefit estimates. We consider that individual-level variation includes explainable variation, for example, due to currently unknown (but feasibly collected) biomarkers, and unexplainable first-order uncertainty due to the inherent differences in outcomes. Both of these sources of variation are distinct from the commonly performed probabilistic analysis (PA; sometimes called probabilistic sensitivity analysis), which explores the impact of second-order, parametric uncertainty on decision making.[20] Note that PA forms the basis of standard VoI calculations.

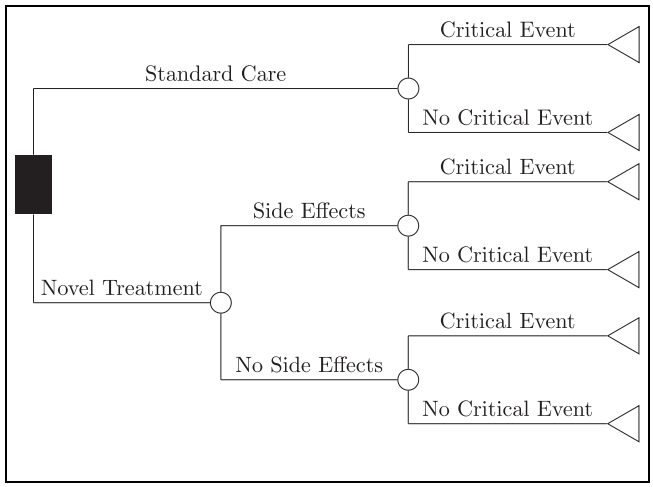

Figure 1 The structure of the decision tree model comparing 2 treatment options for a hypothetical disease that causes a critical event to occur, adapted from Ades et al.[10]

Individualized decision models are generally more complex than cohort models.[21] They often require a higher number of assumptions/parameters as potential relationships between individual-level outcomes and parameters must be considered (e.g., individuals who live longer may have a lower risk of treatment-related adverse events). Detailed, accurate data to inform these individualized models are likely required to ensure the individual-level net benefit is correctly estimated. These data can come from a range of sources but will likely need to be individual patient–level data. In the relatively rare scenario in which clinical trial data are available to accurately define the net benefit,[22] modeling can be minimized and the net benefit estimated directly from these data. The proposed methodology can then be applied to those estimated net benefits.

Individual-Level Decision Models: A Case Study

We developed an individual-level decision model based on a previously published population-level decision tree,[10] depicted in Figure 1. This model calculates the

Table 1 Population-level parameter values for our individual-level decision model comparing treatment to no treatment

Parameter	Description	Value
p_c	Probability of the critical event without treatment	0.4
p_t	Probability of the critical event with treatment	0.2
p_s	Probability of side effects with treatment	0.3
Q_h	Yearly QoL for healthy individuals	1
Q_e	Yearly QoL detriment for individuals who experience the critical event	0.5
Q_s	QoL detriment for individuals who experience side effects	0.1
C_e	Average yearly cost of treating the critical event	£100
V_{C_e}	Variance in the yearly cost of treating the critical event	20
C_s	Average cost of treatment the side effects	£150
V_{C_s}	Variance in the yearly cost of treating the side effects	50
C_t	Cost of treatment	£75
L	Average length of remaining life in the model	3
λ	The willingness-to-pay for a unit of health	20,000

QoL, quality of life.

individual net monetary benefit for 2 treatment options for a hypothetical disease. Under no treatment, modeled individuals are at risk of experiencing a critical event that results in a constant QoL detriment and yearly cost of treatment for the remainder of their life. A treatment is assumed to reduce the risk of this event at the population level. However, this treatment can cause transient side effects, resulting in a short-term QoL detriment and a one-off cost. We assume, as a simplification, that the model parameters, defined in Table 1, are known with certainty. Thus, we do not consider second-order uncertainty, and the optimal treatment, at the population level, is known. If second-order uncertainty is modeled, it can be averaged out before using these methods.

We specify individual-level distributions to generate the individual-specific trajectories, based on these population-level parameters. We model whether an individual experiences the critical event using a Bernoulli distribution, where the probability of experiencing the event is p_c if the individuals receive no treatment and p_t if treated. Similarly, we model whether an individual experiences side effects using a Bernoulli distribution with the probability of experiencing side effects equal to p_s. We then model the individual-level costs of treating the critical event with a mean of C_e and a variance of V_{C_e} using a log-normal distribution with parameters 4.60 and 0.045. Similarly, the cost of treating side effects is assumed to have mean C_s and variance V_{C_s} and follow a log-normal distribution with parameters 5.01 and 0.047. Next, the duration of an individual's life is modeled as an exponential distribution with mean L. We simplify this model by assuming that a patient's length of life is independent of whether they experience the critical event or treatment side effect and that there is no variation in QoL across individuals. Thus, all healthy individuals have the same QoL (Q_h), and any individuals who experience the critical event or side effects will experience the fixed QoL detriments, Q_e and Q_s, respectively. We also fix the treatment cost (C_t) and willingness to pay (λ) across all individuals.

Based on these assumptions, we now present the calculation method for the net monetary benefit for individuals with and without treatment. First, we define 3 health and economic quantities of interest that must be generated to calculate the individual net benefit, irrespective of whether we assume the individual receives treatment or not. We simulate the net monetary benefit for J individuals who do not receive treatment and J individuals who do and define

- $\boldsymbol{I}_e = (I_e^1, I_e^2, \ldots, I_e^J, I_e^{J+1}, \ldots, I_e^{2J})$, a vector of indicators such that $I_e^j = 0$ if individual j does not experience the critical event and $I_e^j = 1$ if they do.
- $\boldsymbol{l} = (l^1, l^2, \ldots, l^J, l^{J+1}, \ldots, l^{2J})$ is a vector representing the length of life for each individual.
- $\boldsymbol{c}_e = (c_e^1, c_e^2, \ldots, c_e^J, c_e^{J+1}, \ldots, c_e^{2J})$ is a vector representing the annual cost of treating the consequences of the critical event for each individual; if $I_e^j = 0$, then $c_e^j = 0$, while if $I_e^j = 1$, then c_e^j is simulated from its log-normal distribution.

We must define 2 additional vectors that capture treatment-related side effects to simulate the individual-level net monetary benefit for J individuals who receive treatment:

- $\boldsymbol{I}_s = (I_s^1, I_s^2, \ldots, I_s^J, I_s^{J+1}, \ldots, I_s^{2J})$ is a vector of indicators of whether the individual experiences side

effects; $I_s^j = 0$ if individual k does not experience the side effect and 1 if they do. Note that if the individual is assumed to not receive treatment, then $I_s^j = 0$ by assumption.

- $\boldsymbol{c}_s = (c_s^1, c_s^2, \ldots, c_s^J, c_s^{J+1}, \ldots, c_s^{2J})$ is a vector representing the one-off cost of treating side effects for each individual. Similar to $\boldsymbol{c}_e$, for individual j, c_s^j is simulated from a log-normal distribution unless $I_s^j = 0$, when c_s^j is set to 0.

Based on these simulations, we can compute the individual-level net monetary benefit for all individuals. We define the first half of the vectors $\boldsymbol{I}_e, \boldsymbol{l}$ and $\boldsymbol{c}_e$ as the outcomes for individuals who do not receive treatment. Thus, the individual-level net benefits without treatment is calculated as, for $j = 1, \ldots, J$,

$$\mathrm{NB}_1^j = \lambda(l^j - l^j I_e^j Q_e) - l^j c_e^j.$$

The first half of this equation, that is, terms multiplied by the willingness to pay (λ), calculates the individual quality-adjusted life-year (QALY) for patient j by subtracting the lifetime QoL detriment due to the critical event from l^j. The second half of the equation calculates the cost of treating the critical event.

The net monetary benefit for individuals receiving treatment is then defined, for $j = J + 1, \ldots, 2J$, as

$$\mathrm{NB}_2^j = \lambda(l^j - l^j I_e^j Q_e - I_s^j Q_s) - l^j c_e^j - c_s^j - C_t.$$

This net benefit calculation also includes the treatment-related side effects that can lead to a QoL detriment and additional costs. Note that, as the side effects only affect individuals in the short term, the QoL detriments and costs are not multiplied by l^j, the individual's length of life.

Once NB_1^j and NB_2^j have been defined and computed, we can calculate the population average net benefit for each treatment option. The optimal treatment if a single strategy was implemented at the population level is the one with the highest average net benefit.[8] In our example, the mean of $\mathbf{NB}_1$ is £4680 and the mean of $\mathbf{NB}_2$ is £5160. Thus, the population-level optimal decision is to treat all individuals.

Maximum VoH

The VoH and EVIC are defined as the expected opportunity loss, across all individuals, incurred by implementing a treatment that is optimal on average, rather than implementing the treatment that is optimal for each individual, separately.[16] At the individual level, the opportunity loss is defined as the difference between the net monetary benefit of the treatment that is optimal for the individual and the net monetary benefit the patient would gain if they were given the treatment that is optimal at the population level.[16]

In the original VoH framework, the individual net benefit is estimated conditional on a set of observed characteristics.[16] However, when generalizing this definition to include unexplained heterogeneity, the individual net benefit is calculated using a model (e.g., our decision tree). More specifically, the individual-level net benefit is calculated as a function of several intermediate simulated quantities, that is, $(\boldsymbol{I}_e, \boldsymbol{I}_s, \boldsymbol{l}, \boldsymbol{c}_e, \boldsymbol{c}_s)$. We will call these intermediate quantities *individual-level outcomes* and denote them $\mathcal{O}$. Thus, the decision model computes the individual-level net benefit for each of D treatment options $d = 1, \ldots, D$ as a function of our individual-level outcomes; $\mathrm{NB}_d(\mathcal{O})$. Using this notation, the maximum VoH (MVoH) as

$$\mathrm{MVoH} = \mathrm{E}\left[\max_d \{\mathrm{NB}_d(\mathcal{O}) - \mathrm{NB}_{d^*}(\mathcal{O})\}\right],$$

where d^* is the optimal treatment at the population level, for example, $d^* = 2$. While this formula uses the same definition as the VoH and EVIC frameworks, we have renamed it to highlight that the heterogeneity in the net benefit includes "unexplained" heterogeneity that derives from the simulated individual-level variation in the outcomes $\mathcal{O}$, whereas VoH and EVIC are concerned only with "explained" heterogeneity that arises from known individual-level preferences or characteristics such as age and gender.

Table 2 visualizes how to estimate the MVoH, based on simulated values for the net monetary benefit for our 2 treatment options. First, the optimal treatment at the population level is found by calculating the average net benefit for each treatment. An individualized treatment decision would treat each individual with the treatment that maximizes their net monetary benefit; for example, individual 1 would not receive treatment while individual 3 would be treated. The opportunity loss of the population-level decision, calculated in column 3, is the difference between the net monetary benefit of the individualized optimal treatment and the net monetary benefit of the population level optimal treatment. Finally, the MVoH is estimated as the average opportunity loss across all individuals.

To compute the MVoH, we must determine the net benefit of each treatment for a given individual. This represents what would have happened to a simulated

Table 2 How to Calculate the Value of Heterogeneity[a]

	NB_1	NB_2	Opportunity Loss
Individual 1	*886*	652	234
Individual 2	*3,679*	1,298	2,381
Individual 3	473	*11,054*	0
Individual 4	*6,518*	-319	6,837
Individual 5	1,121	*3,763*	0
Individual 6	*15,095*	4,875	10,220
Individual 7	1,183	*2,308*	0
Individual 8	6,686	*20,706*	0
Individual 9	*2,059*	302	1,757
⋮	⋮	⋮	⋮
Average	4,660	*5,167*	2,143

[a]The individual net monetary benefit for each treatment is tabulated. The opportunity loss of the population-level decision is the difference between the net monetary benefit of the individual-level optimal treatment and the net monetary benefit of the population level optimal treatment.

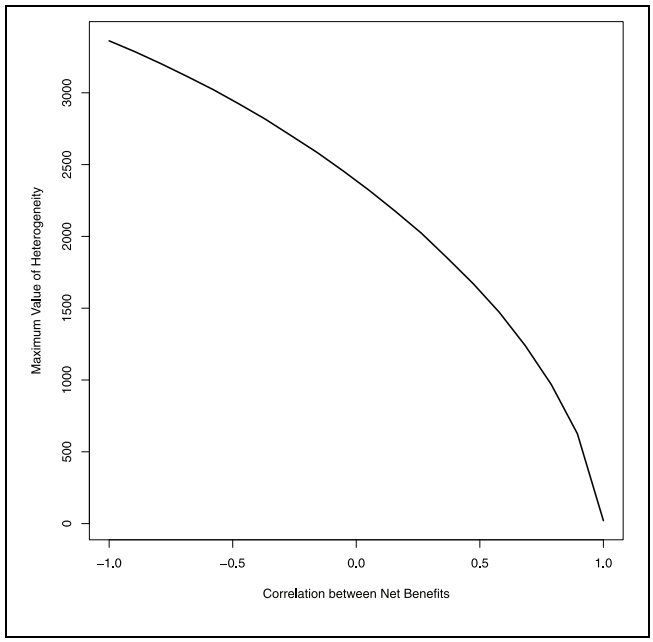

Figure 2 The maximum value of heterogeneity plotted against the correlation between the net benefit across the 2 treatments for the Ades et al.[10] example.

individual under the other treatment options, known as the counterfactual. Espinoza et al. matched patients based on observable characteristics (e.g., age, sex), ensuring that the counterfactual could be computed from individuals with the same characteristics but different treatments.[16] However, the outcomes of a decision model cannot be used to determine the counterfactual based on unknown characteristics. Thus, we must make an explicit assumption about the counterfactual. However, as the magnitude of the opportunity loss is equal to the difference between the net benefit across the different treatment options, the MVoH changes significantly for different assumptions about the counterfactual. Thus, we suggest that these assumptions should be parameterized and the MVoH calculated across a range of scenarios, as we discuss below.

Defining the Counterfactual

The counterfactual can be defined by modifying the correlation, denoted ρ, between the individual net benefits, across the treatment options. Positive values of ρ indicate that some individuals will do well in terms of net monetary benefit, irrespective of treatment option, leading to small opportunity loss and low MVoH. Conversely, negative ρ indicates that individuals would perform well on one treatment option and poorly on another, leading to high opportunity loss and MVoH. Thus, each value for ρ makes a different assumption about the individual-level net benefit across the treatment options, resulting in different counterfactual and different estimates for the MVoH. We will denote the correlation-specific MVoH as MVoH_ρ.

In practice, MVoH_ρ can be estimated from the simulated values of the individual-level net benefit by reordering these simulations so they have a given correlation. Several algorithms are available to reorder simulated values so the resulting data set has a given correlation structure.[23,24] These are implemented in the R package SimJoint,[25] ensuring that the MVoH_ρ can be estimated in standard software.

Figure 2 plots the MVoH analysis for our example. We used $J = 1000$ simulations for individual-level net benefits, $\mathbf{NB}_1$ and $\mathbf{NB}_2$, and computed the MVoH across a range of different values for the Spearman rank correlation ρ^i. The MVoH varies between £20 to £3,300. As ρ decreases, the assumption is that the difference in net benefit between each treatment option increases. This implies that the opportunity loss of selecting the nonoptimal treatment for an individual increases, leading to a larger MVoH.

MVoH is interpreted as the upper bound on the value that could be generated from personalizing treatment decisions. Thus, if the MVoH is "high," then there may be value in developing new personalized medicine strategies. However, the MVoH is calculated conditional on

the correlation between the net benefits ρ, the true value of which is unknown, meaning that the "true" MVoH is not known. Furthermore, the MVoH evaluates that value that could be obtained by explaining all individual-level variation, including all first-order uncertainty, which can never be achieved. Thus, MVoH can be used only to rule out the possibility of increasing value by developing a personalized treatment strategy if MVoH is low for all values of ρ.

The Value of Perfect Outcome Prediction

If the MVoH analysis indicates that there could be substantial value in exploiting the unexplained heterogeneity, it is important to determine what research would allow the development of valuable personalized treatment strategies. In our framework, the heterogeneity in the net benefit comes from the simulated individual-level outcomes $\mathcal{O}$, which are used to compute the net benefits $\mathrm{NB}_d(\mathcal{O})$, $d = 1, \ldots, D$. In the following section, we generalize the parameter-specific EVIC[2] to develop a measure to ascertain which of these outcomes should be predicted to generate value by personalizing treatment. For example, can value be generated by personalizing treatments based on who will experience side effects? Or will personalizing treatment based on knowledge about who experiences the critical event generate greater value?

This prioritization step assumes that, for a set of outcomes $\boldsymbol{o} \subset \mathcal{O}$, the exact value of these outcomes could be predicted for each individual, that is, we know exactly which individuals will experience the critical event with and without treatment. If the exact value of the outcomes $\boldsymbol{o}$ are known for all individuals, we have perfect knowledge of $\boldsymbol{o}$. We can then assess the value of perfect knowledge of $\boldsymbol{o}$, through a measure we have called the value of perfect outcome prediction (VPOP). Note that outcomes are rarely perfectly predictable, as some first-order uncertainty will remain; thus, it is not possible to achieve the VPOP through a realistic personalized treatment strategy. Nevertheless, by calculating the VPOP for different sets of outcomes, we can prioritize research that aims to develop a method to predict outcomes with a high VPOP. Note that the VPOP calculates the value of learning patient outcome (e.g., whether the individual experiences an adverse event), whereas the parameter-specific EVIC is concerned with the value of learning a model parameter (e.g., a QoL weight).

To define the VPOP for a set of outcomes $\boldsymbol{o}$, we denote the remaining outcomes used to compute the individual-level net benefit $\mathcal{O}^-$, that is, $\mathcal{O} = (\boldsymbol{o}, \mathcal{O}^-)$. If the value of $\boldsymbol{o}$ is known, then the individual treatment decision will be made by maximizing the expected net benefit, where expectation is taken over the remaining heterogeneity induced by the unknown outcomes $\mathcal{O}^-$. We then define the VPOP for a set of outcomes $\boldsymbol{o}$ as

$$\mathrm{VPOP}_{\rho_{\boldsymbol{o}}} = \mathrm{E}_{\boldsymbol{o}}\left[\max_d \left\{ \mathrm{E}_{\mathcal{O}^-|\boldsymbol{o}}[\mathrm{NB}_d(\mathcal{O}^-, \boldsymbol{o})] - \mathrm{NB}_{d^*}(\mathcal{O}) \right\} | \rho_{\boldsymbol{o}} \right], \quad (1)$$

where $E_{\mathcal{O}^-|\boldsymbol{o}}[\cdot]]$ is the conditional expectation, implying that perfect information on $\boldsymbol{o}$ may also inform $\mathcal{O}$ if these outcomes are correlated. Similar to the MVoH, the VPOP for a set of outcomes $\boldsymbol{o}$ can be computed only if we make assumptions about the counterfactual so the opportunity loss can be computed. Thus, the VPOP must be recomputed across different assumptions about the correlation between the expected individual-level net benefit across the proposed treatment options, denoted $\rho_{\boldsymbol{o}}$.

The range of possible values for $\rho_{\boldsymbol{o}}$ and its definition will depend on the outcomes included in $\boldsymbol{o}$. If $\boldsymbol{o}$ influences the net benefit for only one of the treatment options, for example, $\boldsymbol{o}$ is whether an individual experienced side effects ($\boldsymbol{I}_s$), then the VPOP does not need to be calculated across different values for $\rho_{\boldsymbol{o}}$, as the expected individual-level net benefit for all other treatments, conditional on $\boldsymbol{o}$, is a single value. Conversely, if $\boldsymbol{o}$ is a continuous outcome used to define the net benefit for all treatment options (e.g., the length of time each individual will live), then $\rho_{\boldsymbol{o}}$ will take a wide range of values, similar to the MVoH analysis.

To compute the VPOP from Eq. 1, we must estimate the expected individual net monetary benefit conditional on the values of $\boldsymbol{o}$—the outcomes we are assuming can be perfectly predicted. The calculation required to achieve this varies depending on whether $\boldsymbol{o}$ contains continuous or only discrete variables. If $\boldsymbol{o}$ includes only discrete variables, then the expected net monetary benefit is estimated by the sample mean of the simulated individual-level net monetary benefit values for all individuals with the same outcome values. Conversely, if $\boldsymbol{o}$ contains continuous variables, then a method based on regression, adapted from Strong et al.,[26] can be used to estimate the conditional monetary net benefit (cf. the supplementary material).

Directing Research with the VPOP

For our example, Figure 3 displays the VPOP for the 5 clinical and economic individual outcomes; the duration of an individual's life ($\boldsymbol{l}$), whether they experience the critical event ($\boldsymbol{I}_c$), treatment side effects ($\boldsymbol{I}_s$), and the cost of treating the critical event ($\boldsymbol{c}_c$) or side effects ($\boldsymbol{c}_s$). The

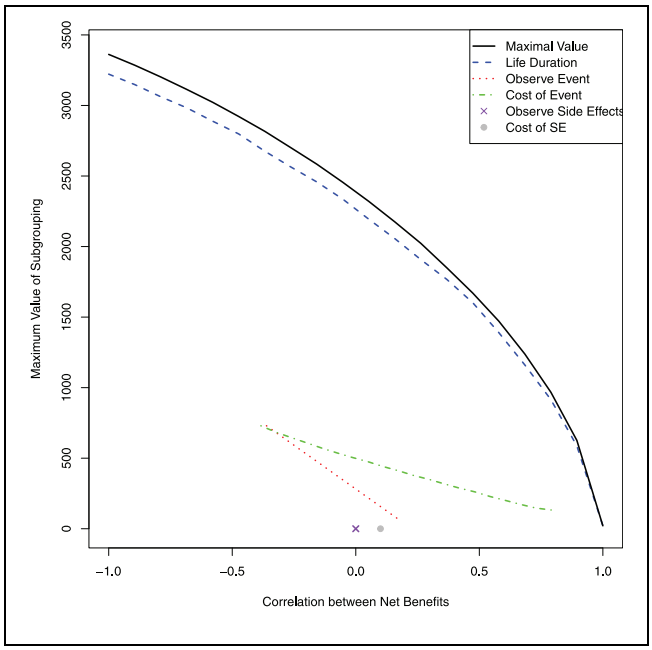

Figure 3 The value of perfect outcome prediction (VPOP) for all outcomes in the modified Ades et al.10 example and the MVoH for comparison (black line). The VPOP for the duration of an individual's life (l) is represented by the blue dashed line. The VPOP for whether an individual experiences the critical event (Ic) is represented by the green dashed and dotted line. The VPOP for the cost of treating the critical event (cc) is represented by the red dotted line. Finally, the VPOP for whether an individual experiences side effects (Is) is represented by a purple cross and the VPOP for the cost of side effects (cs) is represented by a gray dot.

exact method used to compute the VPOP for each these outcomes is presented in the supplementary material.

Figure 3 demonstrates that, for most values of ρ_l, significant value could be generated from a personalized treatment strategy if we could predict life expectancy. The intuition behind this result is that for $\rho_l < 1$, there are some individuals who would have a longer life on one treatment than the other. If we were able to perfectly predict which treatment would allow the individual to live longer, then this should be the treatment they are given. However, by definition in this example, the treatment is not life extending, meaning that we do not expect substantial differences in life expectancy for an individual based on the treatment to which they are assigned. Thus, $\rho_l \approx 1$, and these high values for the VPOP cannot be realized. This argument indicates how the results from a VPOP analysis must be combined with subject knowledge to assess 1) if the correlations between the conditional net benefits are plausible and 2) whether it is practically and/or ethically possible to predict this outcome.

Figure 3 also indicates that predicting whether an individual would experience side effects would not generate value. This is because the risk reduction for the critical event is substantial. Thus, even if the individual would experience side effects, they should still be offered the treatment (cf. the supplementary material). On the other hand, it may be possible to generate value from a personalized treatment strategy, if we knew which individuals would experience the critical event. In this setting, we would be able to identify individuals who do not require treatment (i.e., those who would not experience the critical event even if they receive treatment), and they can avoid being exposed to a risk of harmful side effects.

The Value of Subgroups

Following the VPOP analysis, we now assess whether a test to identify individuals who may experience the critical event could generate sufficient value to encourage its development. In this section, we use scenario analyses to identify what the properties of the test would have to be to generate value from personalizing treatment decisions. We assume that this test will imperfectly predict the outcomes $\boldsymbol{o}$, that is, it will identify subgroups of individuals who may benefit from an alternative treatment. Our scenarios will consider the value of identifying subgroups with different properties. The specification of the relevant properties will vary depending on the exact outcome of interest, the proposed test, and previous information (e.g., from animal models) that could inform the likely outcomes in the different subgroups.

To calculate this value, we assume that the test will identify n_S subgroups. We then generalize Espinoza et al.'s definition of the VoH[16] by denoting the expected net monetary benefit of treatment d for individuals assigned to subgroup s as NB_d^s for $s = 1, \ldots, n_S$ and defining the value of subgroups (VoS) as

$$\mathrm{VoS} = \sum_{s}^{n_S} \omega_s \max_d \mathrm{NB}_d^s - \mathrm{NB}_{d^*}, \qquad (2)$$

where ω_s is the size of subgroup s. Thus, the VoS is a weighted average of the opportunity loss that would be accrued by treating all individuals with the same treatment, rather than the subgroup-specific optimal treatment. We now use our case study to explore the specifics of how the VoS might be calculated in practice.

Value of Subgrouping

We identified from our VPOP analysis that predicting who would experience the critical event with and without

treatment has the potential to generate value. Thus, we define the subgroups in our scenario analysis in terms of the probability of experiencing the critical event. We assume that a test would identify 2 subgroups ($n_S = 2$) where the probability of experiencing the critical event is lower for the individuals in subgroup 1. We also assume that the risk reduction due to treatment is the same across the 2 subgroups. As the critical event has a significant QALY detriment, individuals in subgroup 2 will have a lower net monetary benefit with and without treatment compared with individuals in subgroup 1, implying a positive correlation between the net benefit across the 2 treatments. Thus, the higher values indicated in the VPOP analysis (Figure 3) cannot be achieved with this subgroup structure, as they require a negative correlation between the net benefits across the 2 treatment options.

To define the probability of the critical event in subgroup 1 and subgroup 2 without treatment, p_c^1 and p_c^2, respectively, ($p_c^1 < p_c^2$), we make the assumption

$$p_c = \omega_1 p_c^1 + \omega_2 p_c^2,$$

where $\omega_1 = 1 - \omega_2$. Furthermore, we assume that the reduced probability if the critical event in subgroup 1, p_c^1, is reduced by a multiplicative factor k from the population-level probability of a critical event, for example, $p_c^1 = kp_c$. These 2 assumptions imply that $p_c^2 = \frac{1-k\omega_1}{\omega_2} p_c$. Finally, as the relative risk reduction provided by treatment is assumed to be the same across the 2 groups, the probability of the critical event with treatment in the 2 subgroups is $p_t^s = \frac{p_t}{p_c} p_c^s$, $s = 1, 2$. From these 3 assumptions, we have entirely defined p_c^1, p_c^2, p_t^1 and p_t^2 in terms of the reduction factor k and the size of subgroup 1, ω_1.

Thus, our scenario analyses vary the values of k and ω_1 and then compute the expected net monetary benefit within each subgroup for each treatment. From this, we can use Eq. 2 to compute the VoS. In our example, the net monetary benefits for individuals in each subgroup were calculated by generating $\boldsymbol{I}_c$ with updated values of p_c and p_t. The average net benefit can then be calculated from these simulations and used to estimate the expected net monetary benefit. In complex individual-level decision models or when a large number of scenarios are considered, this analysis will become highly computationally complex.

Figure 4 displays the VoS across different values for the risk reduction in the reduced-risk subgroup, k, and the size of subgroup 1, ω_1, using a heat map. The lighter the color, the higher the VoS, and the more value that could be generated if subgroups with the given values of k and ω_1 were identified. Darker gray indicates that the subgroup structure has no value, and black indicates that the given combination of k and ω_1 are incompatible with the population-level value of p_c. Identifying 2 subgroups in which one has a reduced probability of the critical event will generate value only if more than 40% of the population is in the "reduced risk" subgroup. Furthermore, identifying these subgroups will be valuable only if there is substantial risk reduction. The highest VoS is approximately £100 per individual, achieved if the reduced risk subgroup contains $\sim 60\%$ of the population and have no risk of the critical event.

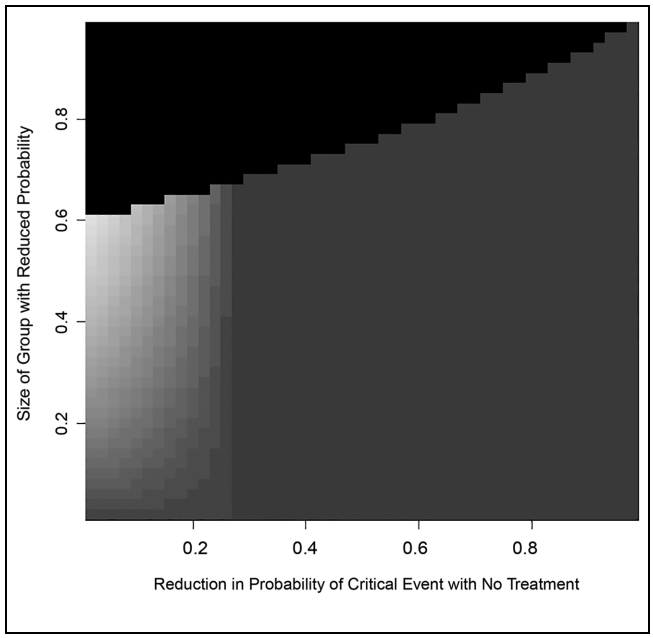

Figure 4 A heat map displaying the value of subgroups for different values of the proportion reduction in the probability of a critical event (k) and the size of the subgroup with the reduced probability of a critical event (ω_1). Darker gray indicates a low value for the subgroups and lighter gray/white indicates a high value. The black section indicates combinations of ω_1 and k that would result in a probability of the critical event in one of the subgroups that is greater than 1.

Discussion

We aimed to develop a method to support research prioritization in the discovery of new personalized treatment strategies, given the limits on research funding. To achieve this, we developed 3 measures that generalize previous work on the VoH[16] to settings in which the source of heterogeneity is unknown. Although research prioritization can also include other considerations such as improving equity in research (e.g., focusing on

underrepresented genders or races), these methods can form part of a comprehensive research prioritization framework in personalized medicine that improves the value of research.

We begin by developing the MVoH, which explores whether value could be generated by explaining individual-level heterogeneity. We then develop the VPOP, which can be used, alongside clinical expertise, to indicate which individual-level outcomes, if they could be predicted, would allow us to personalize treatment and generate value. This measure can prioritize the development of prediction algorithms for specific outcomes. Finally, we demonstrate how scenario analyses can be used to explore the VoS and determine the characteristics that these subgroups need to have to generate value from a potential strategy that offers different treatments to each subgroup. This measure could direct research, as it highlights the features required from a specific test/prediction algorithm to generate value. However, it is more complex to conceptualize and compute as the subgroup structure must be designed separately for each model and the individual decision model must be rerun for each subgroup.

There are several limitations to the currently proposed measures and their graphical representations that should be addressed in future work. First, this article calculates the MVoH and VPOP by varying the correlation between the net benefit to define the counterfactual and then plotting these values against correlation. In decision models with more than 2 treatment options, this method could be extended to estimate the MVoH and VPOP for different correlation matrices. However, the graphical representation used in this article is possible only when the correlation is univariate. Thus, future work should develop alternative graphical presentations for decision models with more than 2 potential treatment options.

Second, we have developed the MVoH, VPOP, and VoS measures based on a deterministic decision model that ignores second-order parametric uncertainty. In general, decision models include a PA to assesses the impact of uncertainty in the model parameters on the decision-making process.[27] Thus, these measures will need to be extended to incorporate parametric uncertainty alongside individual-level heterogeneity, similar to Espinoza et al.[16] Ideally, these extensions would allow researchers to rank the relative importance of reducing uncertainty in key model parameters and developing mechanisms by which we could personalize care.

Third, the proposed VoS is relatively complex to conceive and may be computationally expensive. The relevant scenario analyses will change across different decision models, with the conclusions heavily dependant on the chosen scenarios. There are also challenges associated with presenting the VoS analysis if more than 2 quantities are used to define the subgroups. Thus, future research should focus on conceptualizing and presenting the VoS analysis for different outcome types. Methods will also be required to tackle the computational challenges of this analysis, potentially adapting methods that were developed the reduce the computational cost of VoI analyses,[26,28–32] so they can efficiently compute the net monetary benefit for different subgroup specifications. The VoS analysis will also be most useful when informed by clinical expertise about what subgroups may be potentially available. This will require significant effort to translate clinical expertise into relevant scenarios. If clinical expertise were available, then a VoS analysis could focus on assessing whether realizable subgroup structures could generate value from personalizing treatment. This may reduce the range of scenarios to be considered.

In addition to estimating VoS, research prioritization requires 2 additional elements. Currently available evidence (e.g., from biological models) should be used to consider whether the valuable subgroups are potentially feasible. Second, we should compare the value of subgroups to all the costs of bringing the test to identify the reduced risk subgroup into practice. This could include preclinical and clinical investigation of the test, purchase of equipment to manufacture and analyze the test, obtaining market authorization, and the cost of widespread implementation of the test. These costs could be estimated from the real-world development, validation, and implementation of similar testing procedures.

Finally, these methods will provide accurate research prioritization only if the decision model realistically describes the current observed heterogeneity in the population. This is similar to standard VoI methods in which researchers must accurately capture all second-order uncertainty in their model to ensure correct research prioritization.[33] However, individual-level models are complex to develop, and relevant data may be lacking.[21] Thus, future research should focus on addressing these issues and implementing these measures in a realistic individual-level decision model. However, through the initial presentation of these measures, we have laid the foundation of a framework that will allow researchers to undertake research prioritization for novel personalized treatment strategies.

Acknowledgments

The authors would like to thank David Glynn for his comments on an earlier draft of this article and 3 anonymous

reviewers for their insightful comments. Early ideas included in this article were presented at SMDM2019, Portland, Oregon, in a presentation titled "The Maximum Value of Heterogeneity for Unknown Subgroups."

ORCID iD

Anna Heath https://orcid.org/0000-0002-7263-4251

Supplemental Material

Supplementary material for this article is available on the *Medical Decision Making* website at http://journals.sagepub.com/home/mdm.

Note

i. Difference definitions for the correlation (e.g., a Pearson correlation) provide an alternative relationship between the MVoH and ρ but has limited effect on the range of MVoH values.

References

1. Claxton K. The irrelevance of inference: a decision-making approach to stochastic evaluation of health care technologies. *J Health Econ*. 1999;8:342–64.
2. Basu A, Meltzer D. Value of information on preference heterogeneity and individualized care. *Med Decis Making*. 2007;27(2):112–27.
3. Sculpher M. Subgroups and heterogeneity in cost-effectiveness analysis. *Pharmacoeconomics*. 2008;26(9):799–806.
4. Coyle D, Buxton M, O'Brien B. Stratified cost-effectiveness analysis: a framework for establishing efficient limited use criteria. *Health Econ*. 2003;12(5):421–7.
5. Minelli C, Baio G. Value of information: a tool to improve research prioritization and reduce waste. *PLoS Med*. 2015;12(9):e1001882.
6. Schlaifer R. *Probability and Statistics for Business Decisions*. New York: McGraw-Hill; 1959.
7. Raiffa H, Schlaifer H. *Applied Statistical Decision Theory*. Boston: Harvard University Press; 1961.
8. Claxton K. Bayesian approaches to the value of information: implications for the regulation of new pharmaceutical. *Health Econ*. 1999;8:269–74.
9. Claxton K, Neumann P, Araki S, Weinstein M. Bayesian value-of-information analysis. *Int J Technol Assess Health Care*. 2001;17:38–55.
10. Ades A, Lu G, Claxton K. Expected value of sample information calculations in medical decision modeling. *Med Decis Making*. 2004;24:207–27.
11. Felli J, Hazen G. Sensitivity analysis and the expected value of perfect information. *Med Decis Making*. 1998;18:95–109.
12. Heath A, Hunink M, Krijkamp E, Pechlivanoglou P. Prioritisation and design of clinical trials. *Eur J Epidemiol*. 2021;36:1111–21.
13. Stinnett A, Mullahy J. Net health benefits a new framework for the analysis of uncertainty in cost-effectiveness analysis. *Med Decis Making*. 1998;18(2):S68–80.
14. Willan A, Eckermann S. Optimal clinical trial design using value of information methods with imperfect implementation. *Health Econ*. 2010;19(5):549–61.
15. Fenwick E, Claxton K, Sculpher M. The value of implementation and the value of information: combined and uneven development. *Med Decis Making*. 2008;28(1):21–32.
16. Espinoza M, Manca A, Claxton K, Sculpher M. The value of heterogeneity for cost-effectiveness subgroup analysis: conceptual framework and application. *Med Decis Making*. 2014;34(8):951–64.
17. Agusti A. The path to personalised medicine in copd. *Thorax*. 2014;69(9):857–64.
18. Jackson S, Chester J. Personalised cancer medicine. *Int J Cancer*. 2015;137(2):262–6.
19. Isaacs J, Ferraccioli G. The need for personalised medicine for rheumatoid arthritis. *Ann Rheum Dis*. 2011;70(1):4–7.
20. Briggs A, Weinstein M, Fenwick E, Karnon J, Sculpher M, Paltiel A; ISPOR-SMDM Modeling Good Research Practices Task Force, et al. Model parameter estimation and uncertainty: a report of the ISPOR-SMDM Modeling Good Research Practices Task Force-6. *Value Health*. 2012;15(6):835–42.
21. Krijkamp E, Alarid-Escudero F, Enns E, Jalal H, Hunink M, Pechlivanoglou P. Microsimulation modeling for health decision sciences using r: a tutorial. *Med Decis Making*. 2018;38(3):400–22.
22. Sculpher MJ, Claxton K, Drummond M, McCabe C. Whither trial-based economic evaluation for health care decision making? *Health Econ*. 2006;15(7):677–87.
23. Iman R, Conover W. A distribution-free approach to inducing rank correlation among input variables. *Commun Stat Simulat Comput*. 1982;11(3):311–34.
24. Ruscio J, Kaczetow W. Simulating multivariate nonnormal data using an iterative algorithm. *Multivariate Behav Res*. 2008;43(3):355–81.
25. Liu CW. *SimJoint: Simulate Joint Distribution*. R package version 0.3.7. 2020.
26. Strong M, Oakley J, Brennan A. Estimating multiparameter partial expected value of perfect information from a probabilistic sensitivity analysis sample: a nonparametric regression approach. *Med Decis Making*. 2014;34(3):311–26.
27. Claxton K, Sculpher M, McCabe C, et al. Probabilistic sensitivity analysis for NICE technology assessment: not an optional extra. *Health Econ*. 2005;14:339–47.
28. Strong M, Oakley J, Brennan A, Breeze P. Estimating the expected value of sample information using the probabilistic sensitivity analysis sample: a fast nonparametric regression-based method. *Med Decis Making*. 2015;35(5):570–83.
29. Heath A, Manolopoulou I, Baio G. Estimating the expected value of partial perfect information in health economic evaluations using integrated nested laplace approximation. *Stat Med*. 2016;35(23):4264–80.

30. Menzies N. An efficient estimator for the expected value of sample information. *Med Decis Making*. 2016;36(3): 308–20.
31. Jalal H, Alarid-Escudero F. A Gaussian approximation approach for value of information analysis. *Med Decis Making*. 2017;38(3):174–88.
32. Heath A. and Baio G. Calculating the expected value of sample information using efficient nested Monte Carlo: a tutorial. *Value in Health*, 2018;21(11):1299-1304.
33. Koffijberg H, Rothery C, Chalkidou K, Grutters J. Value of information choices that influence estimates: a systematic review of prevailing considerations. *Med Decis Making*. 2018;38(7):888–900.