Biometrical Journal

WILEY-VCH

Biometrical Journal

RESEARCH ARTICLE OPEN ACCESS

Risk-Based Decision Making: Estimands for Sequential Prediction Under Interventions

Kim Luijken[1] | Paweł Morzywołek[2,3] | Wouter van Amsterdam[1] | Giovanni Cinà[4,5,6] | Jeroen Hoogland[7] | Ruth Keogh[8] | Jesse H. Krijthe[9] | Sara Magliacane[10] | Thijs van Ommen[11] | Niels Peek[12] | Hein Putter[13] | Maarten van Smeden[1] | Matthew Sperrin[12] | Junfeng Wang[14] | Daniala L. Weir[14] | Vanessa Didelez[15,16] | Nan van Geloven[13]

[1]Department of Epidemiology, Julius Center for Health Sciences and Primary Care, University Medical Center Utrecht, Utrecht, the Netherlands | [2]Department of Applied Mathematics, Computer Science and Statistics, Ghent University, Ghent, Belgium | [3]Department of Statistics, University of Washington, Seattle, Washington, USA | [4]Department of Medical Informatics, Amsterdam University Medical Centers, Amsterdam, the Netherlands | [5]Institute for Logic, Language and Computation, University of Amsterdam, Amsterdam, the Netherlands | [6]Pacmed, Amsterdam, the Netherlands | [7]Department of Epidemiology and Data Science, Amsterdam University Medical Centers, Amsterdam, the Netherlands | [8]Department of Medical Statistics, London School of Hygiene & Tropical Medicine, London, UK | [9]Pattern Recognition and Bio-Informatics Group, EEMCS, Delft University of Technology, Delft, the Netherlands | [10]Amsterdam Machine Learning Lab, University of Amsterdam, Amsterdam, the Netherlands | [11]Department of Information and Computing Sciences, Utrecht University, Utrecht, the Netherlands | [12]Division of Informatics, Imaging and Data Science, Faculty of Biology, Medicine and Health, University of Manchester, Manchester Academic Health Science Centre, Manchester, UK | [13]Department of Biomedical Data Sciences, Leiden University Medical Center, Leiden, the Netherlands | [14]Division of Pharmacoepidemiology and Clinical Pharmacology, Department of Pharmaceutical Sciences, Utrecht University, Utrecht, the Netherlands | [15]Department of Biometry and Data Management, Leibniz Institute for Prevention Research and Epidemiology - BIPS, Bremen, Germany | [16]Faculty of Mathematics/Computer Science, University of Bremen, Bremen, Germany

Correspondence: Nan van Geloven (n.van_geloven@lumc.nl)

Received: 29 November 2023 | **Revised:** 16 June 2024 | **Accepted:** 24 October 2024

Funding: The authors received no specific funding for this work.

Keywords: counterfactual prediction | estimand | prediction model | prediction under interventions

ABSTRACT

Prediction models are used among others to inform medical decisions on interventions. Typically, individuals with high risks of adverse outcomes are advised to undergo an intervention while those at low risk are advised to refrain from it. Standard prediction models do not always provide risks that are relevant to inform such decisions: for example, an individual may be estimated to be at low risk because similar individuals in the past received an intervention which lowered their risk. Therefore, prediction models supporting decisions should target risks belonging to defined intervention strategies. Previous works on prediction under interventions assumed that the prediction model was used only at one time point to make an intervention decision. In clinical practice, intervention decisions are rarely made only once: they might be repeated, deferred, and reevaluated. This requires estimated risks under interventions that can be reconsidered at several potential decision moments. In the current work, we highlight key considerations for formulating estimands in sequential prediction under interventions that can inform such intervention decisions. We illustrate these considerations by giving examples of estimands for a case study about choosing between vaginal delivery and cesarean section for women giving birth. Our formalization of prediction tasks in a sequential, causal, and estimand context provides guidance for future studies to ensure that the right question is answered and appropriate causal estimation approaches are chosen to develop sequential prediction models that can inform intervention decisions.

1 | Introduction

To enhance health and health care, there is a need for smart decision support tools that can improve medical decision making by providing care professionals and patients with person-specific information. Clinical prediction models estimate risks of (future) outcomes conditional on patient characteristics and thus have the potential to provide such information (Moons et al. 2015; Moons, Royston, et al. 2009; Riley et al. 2019; Steyerberg 2009). Yet, standard prediction models are generally not suitable for supporting intervention decisions. Prediction models that are intended to inform decisions about interventions need to answer questions like "what is the risk of outcome Y under intervention option a conditional on an individual's characteristics $\mathbf{X}$?" Answering such "what if" questions requires embedding in causal reasoning (Coston et al. 2020; Dickerman et al. 2022; Dickerman and Hernán 2020; Lin et al. 2021; Sperrin et al. 2018, 2019; van Geloven et al. 2020, forthcoming). We refer to this task as "prediction under (hypothetical) interventions."[1]

A number of studies have emphasized the need for prediction under interventions and clarified how to estimate risks under interventions (Coston et al. 2020; Dickerman and Hernán 2020; Dickerman et al. 2022; Lin et al. 2021; Sperrin et al. 2018, 2019; van Geloven et al. 2020). The focus of these studies was on intervention decisions that were made at a single time point, but in clinical practice intervention decisions are rarely made only once and might often be deferred or reevaluated. To support such sequential decisions, prediction models are needed that can estimate risks under interventions at any potential decision moment. By sequentially providing estimated outcome risks under intervention options conditional on individual characteristics, the model functions as an *assistive* decision support tool (Reilly and Evans 2006). Methods for sequential predictions are well established in a standard prediction context, but not yet for interventional predictions (Rizopoulos 2012; van Houwelingen and Putter 2011). The current work clarifies how to formulate and formalize sequential predictions under interventions.

The topic of the current work, *sequential prediction under interventions*, is different from the related topic of optimal dynamic intervention regimes. A dynamic intervention regime is a function that takes baseline covariates, covariate history, and intervention history as inputs and returns an intervention decision to be made next (Chakraborty and Moodie 2013; Dawid and Didelez 2010; Moodie, Richardson, and Stephens 2007; Murphy 2003; Robins 2004). Methods exist for deriving optimal dynamic intervention regimes which are optimal with respect to a predefined utility function. Optimal dynamic intervention rules provide advice on what intervention to assign at a moment in time. Methods for deriving optimal dynamic decision rules assume that an evaluation function exists to assess the value of candidate decision rules. In practice, the challenge is often to come up with such a utility function: it will determine everything. Our work can provide utility functions—in particular, those that assess the risk of adverse health outcomes under a given candidate decision rule. Simply providing information on outcome risks under interventions can sometimes be more useful to end users than a suggested intervention decision. This is because patients and care professionals can have different utility functions and their value judgments and view on what is optimal can change over time. An optimal intervention rule requires assignment of an intervention as recommended by the decision rule during all stages in order to arrive at the optimal outcome. Care professionals may be hesitant to consistently follow an algorithm's suggestions for multiple intervention decisions ahead. Instead, having estimated risks of outcomes under intervention options on the table can facilitate a conversation between care professional and patient so that informed decisions can be made (Morzywołek 2023; van Geloven et al. 2020).

The starting point of development of a prediction model under interventions is formulating the risk questions that reflect the desired predictions, that is, formulating the estimands. The current work highlights key considerations for formulating estimands for sequential prediction under interventions. We illustrate these considerations by giving examples of estimands for single-stage and sequential prediction under interventions in a case study about deciding on vaginal delivery or cesarean section for women giving birth. The contribution of this paper is to define prediction tasks in a sequential, causal, and estimand context. A clearly defined estimand is needed to ensure answering the right question and choosing appropriate causal estimation approaches.

2 | How to Formulate Estimands for Prediction Under Interventions

We start with providing general considerations for formulating estimands for prediction under interventions and then highlight aspects that require additional attention in the sequential setting.

An estimand is a precise definition of the target quantity of an analysis. In prediction studies, an estimand captures the conditional outcome risk (or value) of interest that is intended to be used as information about a particular individual. The conditional outcome risk can be interpreted as *individualized* rather than individual, because they are formalized at a group level given a covariate pattern (Hoogland et al. 2021; Knaus, Lechner, and Strittmatter 2021). The term "estimand" is rarely used in prediction modeling studies, but prediction guidelines implicitly provide recommendations on aspects that should be defined to specify the target of a prediction study (Luijken et al. 2019; Moons et al. 2014, 2015; Whittle et al. 2017). By combining these recommendations with roadmaps available from clinical trial literature (European Medicines Agency 2020) and causal inference literature (Goetghebeur et al. 2020; Petersen and van der Laan 2014), we arrive at the following elements that are required for estimands for prediction under interventions:

- Population: A characterization of the target population to whom the prediction model is to be applied, including the care setting;

[1] Prediction under interventions has also been referred to as "counterfactual prediction" (Coston et al. 2020; Dickerman and Hernán 2020; Hernán, Hsu, and Healy 2019). We use the term prediction under interventions instead because at the moment of making the prediction the outcomes are still in the future and hence risks need not be counterfactual (Dawid and Didelez 2010).

- Moment(s) of intended use: The moment(s) at which the prediction model is to be used to inform the intervention decision;
- Intervention options: The interventions considered at the moment(s) of intended use and, if relevant, for how long they are administered;
- Outcome and prediction horizon: The predicted outcome including the time since the moment of prediction at which we consider it. The specified outcome and prediction horizon should represent information that is needed in the discussion among care professionals and patients to make the intervention decision;
- Predictor(s): The measurement(s) used for prediction that characterize the individual, that is, that individualize the predictions. The measurement(s) must be available at the moment(s) of intended use and measurement procedures must correspond to those in the setting of intended use.

2.1 | Considerations for Prediction Under Interventions at a Single Time Point

As an example of a model that provides predictions under interventions at a single time point, an estimand could express the risk of a binary outcome $Y \in \{0,1\}$ at a prediction horizon h, conditional on predictors $\mathbf{X}$ measured at moment of intended use under specified intervention options $A \in \{0,1\}$ in a given population. Introducing notation, we denote regular time points by $k = 0, 1, 2, \ldots, K$. Observed intervention status A_k, predictors $\mathbf{X}_k$, and outcome status Y_k are measured at each time point. We assume the moment of intended use is time point $k = 0$. We let $\underline{a}_0 = (a_0, \ldots, a_K)$ be a specified joint intervention, setting the intervention level a_k for each time point between 0 and K. In case, the levels are "perform the intervention" ($a = 1$) or "do not perform the intervention" ($a = 0$), we can denote by $\underline{a}_0 = \mathbf{0}$ that the intervention is set to 0 at all time points, which then indicates not giving an intervention at the moment of prediction and continuing to do so until time point K; $\underline{a}_0 = \mathbf{1}$ then indicates giving the intervention at the moment of prediction and continuing it until time point K. An intervention option might also refer to a well-defined standard of care. We define the potential outcome $Y_h^{\underline{a}_0}$ as the outcome at prediction horizon h if an individual is assigned intervention sequence $\underline{a}_0$. By fixing the intervention to a certain strategy, an estimand for prediction under intervention at a single time point can be formally defined as

$$\Pr\left[Y_h^{\underline{a}_0} = 1 | \mathbf{X_0}\right],$$

for example, for $\underline{a}_0 = \mathbf{0}$, this would be the conditional risk of outcome Y at prediction horizon h for an individual with predictor values $\mathbf{X_0}$ if they would not initiate the intervention up to the prediction horizon.

To precisely define the intervention options for the estimand, it is helpful to make a distinction between different types of interventions. Point interventions are administered once or for a (very) short duration, like single-dose pharmacological treatments or surgery. Sustained interventions are administered during longer disease episodes and can take the form of a *static* or *dynamic* regime (Chakraborty and Moodie 2013; Sterne et al. 2016). A sustained static regime specifies a fixed sequence of the intervention status for a predefined duration, like $\underline{a}_0$ in the estimand defined above. Or, for instance, "an individual starts the intervention of interest and continues to use it during a year." Sustained static regimes can be seen as joint interventions as they do not take information into account that occurs during the course of time. Dynamic regimes were mentioned in the introduction and propose an intervention decision based on the (time-varying) health state of an individual. For instance, "an individual starts the intervention of interest when marker X drops below a certain value." Point interventions, sustained static intervention regimes, and dynamic intervention regimes can all be used as intervention options in an estimand for prediction under interventions.

2.2 | Considerations for Sequential Prediction Under Interventions

When an intervention decision can be deferred or reevaluated, a prediction model capable of sequentially estimating risks under hypothetical interventions is needed.

We consider a setting in which predictions of a binary outcome are made at regular time points $k = 0, 1, 2, \ldots$. The prediction horizon can either stay fixed for each k, that is, $h_k = h$, or can be defined relative to the moment of prediction, for example, $h_k = k + w$. Let $\underline{a}_k = (a_k, \ldots, a_K)$ be the specified intervention sequence of interest. We let $\overline{\mathbf{X}}_k = (\mathbf{X}_0, \mathbf{X}_1, \mathbf{X}_2, \ldots, \mathbf{X}_k)$ denote the observed predictor history up to k, $\overline{A}_k = (A_0, A_1, A_2, \ldots, A_k)$ denote the observed intervention history up to k, and Y_k denote whether the outcome has occurred before time k. An estimand for sequential prediction under an intervention option can then be defined as

$$\Pr\left[Y_{h_k}^{\underline{a}_k} = 1 | \overline{\mathbf{X}}_k, Y_k = 0, \overline{A}_{k-1}\right],$$

for example, for $\underline{a}_k = \mathbf{0}$ and $\overline{A}_{k-1} = \mathbf{0}$, this estimand reflects the conditional outcome risk at time $k + w$ for an individual with a predictor history $\overline{\mathbf{X}}_k$ who did not experience the outcome or initiate the intervention before time k, if they would not initiate the intervention up to the prediction horizon by deciding on "no intervention" each time the need for intervening is reevaluated.

An estimand for sequential prediction under interventions incorporates updated predictor information. The time until the next prediction moment informs how long the intervention options could be fixed and at what time point the prediction horizon could be set. For instance, if prediction moments follow each other at short intervals, it might be more reasonable to define a short-term prediction horizon and to fix the intervention option to a certain level until that prediction horizon. This is likely less reasonable if predictions under interventions are made infrequently. An overview of key considerations regarding estimands for sequential prediction under interventions is given in Table 1.

TABLE 1 | Considerations to define estimands for sequential prediction under interventions.

Estimand element	Questions that help formulate the estimand element
Population	• To which individuals will the prediction model be applied?
	• In which health care setting will the prediction model be applied?
Moment(s) of intended use	• At which moment(s) is the prediction model (re)consulted to inform the intervention decision?
Intervention options	• Which intervention options are relevant at the moment(s) of making the intervention decision?
	• For how long should the intervention strategy be fixed?
	• Should the duration to fix the intervention option be aligned with the time till next moment of prediction?
Outcome and prediction horizon	• Which outcome(s) are most informative for the intervention decision?
	• What prediction horizon provides important information for the intervention decision: a short-term or long-term horizon?
	• Should the outcome be defined differently because of the specified intervention option(s)?
	• Should the prediction horizon be aligned with the time till next moment of prediction?
Predictor(s)	• Which predictors are predictive of the outcome of interest?
	• Based on which characteristics should the outcome risks be individualized?
	• Which measurements are available at the moment(s) of intended use?

3 | Case study: Mode of Delivery in Women With High-Risk Pregnancies

We use a clinical case study to illustrate the required considerations for formulating estimands for sequential prediction under interventions. We define several example estimands for a prediction model that informs the intervention decision for mode of delivery in women giving birth. We start out with a setting in which the prediction model is only used once to inform the decision. For this single-stage prediction, we formulate four example estimands that vary in their intervention options. The example is then extended to sequential prediction under interventions with three example estimands that vary in intervention options and prediction horizon.

3.1 | Clinical Context

The clinical case study is inspired by a prediction model developed by Schuit and colleagues that was proposed to support medical decisions during labor of high-risk pregnant women (Schuit et al. 2012). Different from the original model, in the current study we define the outcome as a composite of adverse events in the mother and child. To keep our discussion focused, we simplify the intervention options and focus on vaginal delivery and cesarean section only. Our example estimands can be extended to include other intervention options, for example, ways to induce natural vaginal delivery or instrumental vaginal delivery, in a similar manner.

As a brief background on the clinical setting based on Schuit and colleagues (Schuit et al. 2012); some pregnancies are classified as "high-risk" because they are complicated by pre-existing maternal disease or complications during pregnancy. High-risk pregnant women are typically monitored by gynecologists in secondary care and are admitted to hospital to give birth. Birth can take place via different modes of delivery including natural vaginal delivery and cesarean section. Cesarean section can be a way to prevent adverse neonatal outcomes due to fetal distress, but it comes with risks for the mother like increased blood loss, incontinence, and infection (National Institute for Health and Care Excellence 2021). The decision to give birth via cesarean section can be made before start of labor so that a planned cesarean section can be performed at an elected time (National Institute for Health and Care Excellence 2019). When high-risk pregnant women not scheduled for planned cesarean section go into labor, the decision about mode of delivery is reevaluated at start of labor and repeatedly during labor. The decision to perform cesarean section is based on a weighting of benefits and harms of the procedure and preferences of care professionals and the woman giving birth (National Institute for Health and Care Excellence 2023, 2019).

3.2 | Elements Shared Across the Example Estimands

Common elements of the estimands formulated in the case study are:

- Population: Pregnant women at a gestational age of 36 weeks or over with preexisting maternal disease or complications during pregnancy not scheduled for planned cesarean section admitted to the hospital to give birth;

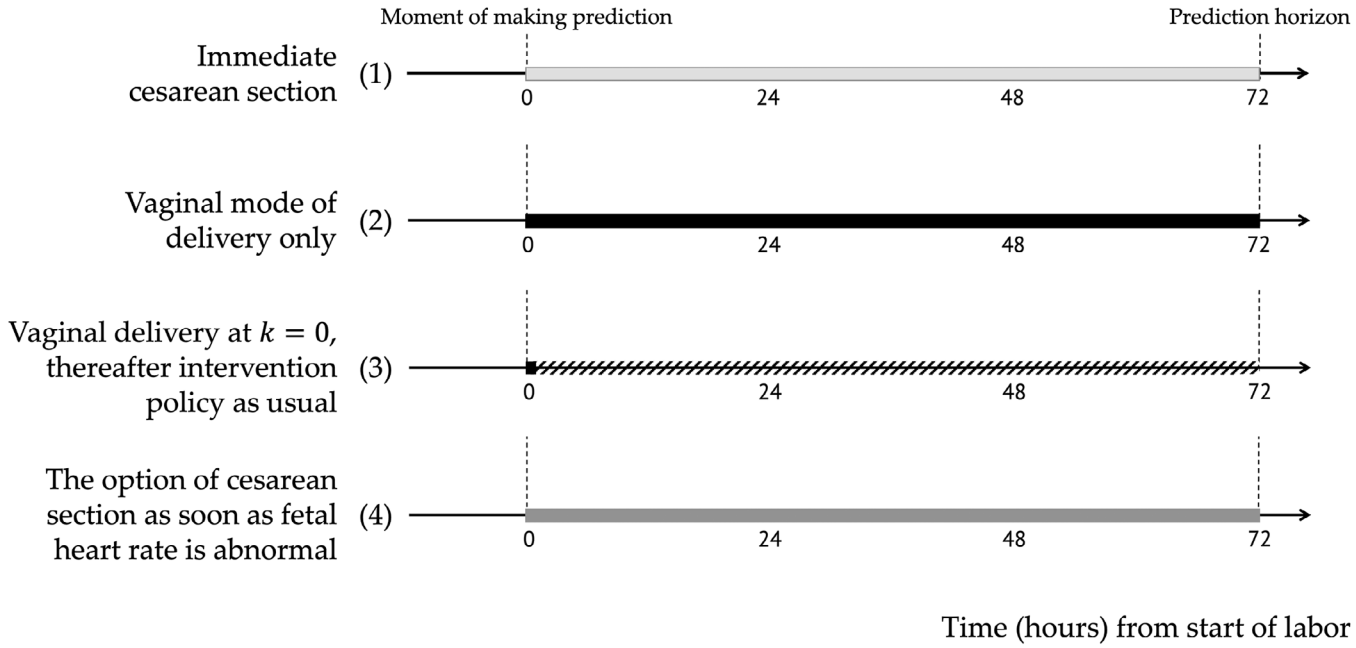

FIGURE 1 | Illustrations of estimands for single-stage prediction under intervention options in the case study. The shaded rectangles indicate the intervention options under which predictions are made: (1) Light gray for cesarean section at time point 0, which is irreversible; (2) Black for vaginal mode of delivery only; (3) Black for setting mode of delivery to vaginal delivery at time point 0, and diagonal stripes for the option of cesarean section as in usual care; (4) Dark gray for the dynamic intervention rule "performing a cesarean section as soon as fetal heart rate is abnormal."

- Moment(s) of intended use: At the start of labor for the single-stage prediction and hourly from start of labor for the sequential prediction;
- Intervention options: Intervention options considered are two modes of delivery. The first is continuing the delivery vaginally and to refrain from cesarean section for a certain period of time. The second is to perform a cesarean section. The duration of fixing the intervention varies across estimands;
- Outcome and prediction horizon: The outcome is defined as a composite of any adverse neonatal or maternal outcomes. Adverse neonatal outcomes include stillbirth, early neonatal death, or requiring admission to a neonatal intensive care unit. Adverse maternal outcomes include maternal death and postpartum admission to intensive care. The prediction horizon varies across example estimands;
- Predictors: Fetal heart rate, dilatation, maternal systolic blood pressure, maternal diastolic blood pressure, maternal age, parity, and history of preterm birth. For the sequential prediction, we additionally use hourly updated information of the time-varying predictors (fetal heart rate, dilatation, maternal systolic blood pressure, and maternal diastolic blood pressure).

3.3 | Single-Stage Prediction Under Interventions

In the single-stage prediction setting, we assume the prediction model is used only at the start of labor to estimate risks of adverse outcomes 72 h later under different intervention options. We vary the defined intervention options across four example estimands (Figure 1). Note that a single-stage prediction assumes a single time point at which predictions are made to inform the intervention decision, but these predictions can apply to a sequence of (hypothetical) future interventions.

3.3.1 | Notation

We introduce some further notation specific to the case study. Let $k = 0, 1, 2, \ldots, 72$ be time points of 1 h apart, where $k = 0$ denotes the start of labor and $h = 72$ denotes the prediction horizon. Let $\underline{a}_0 = (a_0, a_1, a_2, \ldots, a_{71})$ denote an intervention option for the 72 time intervals, where $a_k \in \{0, 1\}$, $a_k = 1$ denotes cesarean section at time k and $a_k = 0$ denotes deciding on vaginal mode of delivery at time k. Because cesarean section is irreversible, a_k stays 1 after it was set to 1 for the first time by the nature of the intervention. We use $\underline{A}_k = (A_k, \ldots, A_{71})$ to denote a vector of observed intervention status between time k and time 72. Let $\mathbf{X_0}$ denote the set of seven predictors available at the start of labor: fetal heart rate, dilatation, maternal systolic blood pressure, maternal diastolic blood pressure, maternal age, parity, and history of preterm birth. Let $Y_{72} \in \{0, 1\}$ denote the outcome status at the prediction horizon, where $Y_{72} = 1$ represents any of the defined adverse neonatal or maternal outcomes having occurred and $Y_{72} = 0$ represents completion of birth in absence of the adverse outcomes. Because all women can be expected to have given birth after 72 h, the chosen prediction horizon implies that the outcome reflects all adverse events.

3.3.2 | Estimand 1: Immediate Cesarean Section

To formulate the outcome risk when cesarean section is performed immediately at the start of labor, we define the

intervention option as setting mode of delivery to cesarean section at $k = 0$, that is, $a_0 = 1$. Then,

$$\Pr\left[Y_{72}^{a_0=1} = 1|\mathbf{X_0}\right] \quad (1)$$

expresses the conditional risk of adverse neonatal or maternal outcomes 72 h after start of labor under immediate cesarean section at start of labor given the predictors at start of labor. As indicated before, because cesarean section is irreversible, Estimand (1) corresponds to $\Pr(Y_{72}^{\underline{a}_0=\mathbf{1}} = 1|\mathbf{X_0})$. Note that predictors after time point 0 cannot be considered, because this would imply that model users would need to input information about the future at moment of intended use.

3.3.3 | Estimand 2: Vaginal Mode of Delivery Only

Another intervention option would be a sustained static intervention regime from the start of labor until the prediction horizon, for example, $\underline{a}_0 = \mathbf{0}$. Then,

$$\Pr\left[Y_{72}^{\underline{a}_0=\mathbf{0}} = 1|\mathbf{X_0}\right] \quad (2)$$

expresses the conditional risk of adverse neonatal or maternal outcomes 72 h after start of labor under vaginal mode of delivery only, given the predictors at start of labor.

3.3.4 | Estimand 3: Deciding on Vaginal Mode of Delivery at Moment of Prediction and Thereafter Intervention Policy "As Usual"

Estimand (2) informs the intervention decision by predicting outcomes under perfect adherence to the intervention. The woman in labor and health care professionals may prefer a vaginal delivery but the intervention strategy "vaginal mode of delivery only" likely does not reflect actual practice. In practice, women and care professionals reevaluate the need for intervening and may decide to change the mode of delivery. Information on the conditional risk of adverse events under practice as usual can be obtained, for example, by setting an intervention option only for the moment at which the prediction is made, without committing to a particular intervention choice or decision rule at future time points. Predicting under usual practice refers to the typical practice and decision rules used in the development/training data.

To estimate such risks, we define the intervention option as setting mode of delivery to vaginal delivery between $k = 0$ and $k = 1$, that is, $a_0 = 0$, and leaving the option of choosing cesarean section later if indicated according to usual intervention assignment during all subsequent times k, that is, $\underline{A}_1$ (Morzywołek 2023, Ch. 5). This implies that a_0 is fixed to a value, whereas $\underline{A}_1$ is a random variable taking on the "natural" values, that is, what would happen under the policy used in the development/training data. Then,

$$\Pr\left[Y_{72}^{a_0=0,\underline{A}_1} = 1|\mathbf{X_0}\right] \quad (3)$$

expresses the conditional risk of adverse neonatal or maternal outcomes 72 h after the start of labor under deciding on vaginal mode of delivery for the first hour after start of labor and the option of cesarean section at subsequent time points according to usual intervention policy, given the predictors at start of labor. The way that the intervention option is specified in Estimand (3) relates closely to the way this is done in the intervention contrast of an intention to treat effect.

The appeal of Estimand (3) is that the intervention option is similar to the way a decision might be made in clinical practice: an intervention option is now chosen, but there is no commitment to a particular intervention option or decision rule at future points in time. The downside of this is that the conditional outcome risks are dependent on the intervention assignment that was observed in the development/training data. This implies that estimated risks will only generalize to settings with similar intervention assignment policies (van Geloven et al. 2020). Predictions under a predefined static sustained intervention like in Estimands (1) and (2) or under a (optimal) dynamic intervention do not rely on this assumption.

3.3.5 | Estimand 4: Performing a Cesarean Section as Soon as Fetal Heart Rate is Abnormal

At the start of labor, a possible decision is to perform a cesarean section as soon as but only if there are signs of fetal heart rate abnormality. More specifically, we define a dynamic intervention rule that cesarean section is performed the first time fetal heart rate is less than 110 beats per minute for 3 min (persistent fetal bradycardia) or higher than 160 beats per minute (fetal tachycardia). This is an example of a basic dynamic intervention rule, in which an intervention is assigned based on the history of a single covariate up to that stage. More generally, dynamic intervention rules can also be based on multiple covariates such as taking the mother's health status into account. A dynamic intervention rule could also be an optimal, as described in the introduction.

Let X_k denote fetal heart rate at time point k and g_k an indicator function that flags abnormal fetal heart rate, where

$$g_k(x_k) = \begin{cases} 1 & \text{if } x_k < 110 \text{ for 3 min in a row or } x_k > 160 \\ 0 & \text{otherwise.} \end{cases}$$

The dynamic intervention sequence d equals 0 until the first time point k where $g_k(X_k) = 1$, and 1 from that point on wards. As the X_k's are random, the sequence d is random as well, we only fix the rule on how to respond to variations in X_k. The estimand of interest can be specified as

$$\Pr\left[Y_{72}^{\underline{a}_0=d} = 1|\mathbf{X}_0\right]. \quad (4)$$

This estimand expresses the conditional risk of adverse neonatal or maternal outcomes under the intervention option to perform cesarean section as soon as the fetal heart rate is abnormal, given the predictors at start of labor.

It deserves clarification as to why we consider a dynamic intervention strategy for the single-stage prediction setting, which may be a confusing classification. The dynamic intervention rule states that based on fetal heart rate during the course of labor, cesarean section can be performed. This means that the need for intervening is reevaluated over time. However, the choice to *commit to this rule* is made only once, at the start of labor. In other words, the action (of starting cesarean section) is reevaluated, the decision is not. To support the decision, we only need a single-stage prediction of the conditional outcome risk under the dynamic intervention rule (Figure 1—(4)). Next, we will consider estimands where the prediction under intervention options is reevaluated over time.

3.4 | Sequential Prediction Under Interventions

Predictions under interventions can be considered several times during the course of labor to inform the decision on mode of delivery. We continue the example by assuming that the prediction model is used at the start of labor and thereafter revisited every hour to update the risks. When defining an estimand that is suitable for sequential prediction under interventions, it is important to consider the time-varying information up to the moment of making the prediction as well as the time axes for the intervention options and outcome after the moment of making the prediction. We discuss these considerations using three examples of relevant estimands.

3.4.1 | Further Notation

We now assume that predictions informing the decision about mode of delivery are made every hour $k = 0, 1, 2, \ldots$. We keep the prediction horizon at 72 h after start of labor. In other settings, the prediction horizon may be defined relative to the moment of making the prediction, as explained in Section 2.2. The predictors $\mathbf{X}_k$ are measured at each time point.[2] Let $\overline{\mathbf{X}}_k = (\mathbf{X}_0, \mathbf{X}_1, \mathbf{X}_2, \ldots, \mathbf{X}_k)$ denote the predictor history up to time point k. We denote whether the woman has given birth at time k by introducing an "at risk" indicator $Z_k \in \{0, 1\}$, where $Z_k = 1$ denotes that the woman is still in labor and no adverse outcome has occurred yet and $Z_k = 0$ denotes that the woman has given birth or an adverse outcome has occurred.

3.4.2 | Estimand 5: Vaginal Mode of Delivery Only, Using Updated Predictor Information

We first extend Estimand (2) to a sequential prediction setting using updated information at the moment of prediction. The starting point is to consider for which individuals the intervention decision still is to be made, that is, which individuals form the "risk set" at time k. In this example, the intervention decision on mode of delivery is to be made in women who are still in labor and who have not experienced an adverse event. Particularly, women who underwent a cesarean section are no longer part of the risk set after administering the intervention. We use $Z_k = 1$ to select the risk set of women who are still in labor and have not experienced an adverse event by time k. In general, one would condition on outcome status ($Y_k = 0$) and intervention history ($\overline{A}_{k-1}$), as is shown in Section 2.2. In the setting we consider, the conditioning on the at risk indicator $Z_k = 1$ implies $Y_k = 0$ and $\overline{A}_{k-1} = \mathbf{0}$.

Predictor information up to the moment of making the prediction is $\overline{\mathbf{X}}_k$. We define the intervention option as setting mode of delivery to vaginal delivery from time k onward for the remaining time of labor, that is, $\underline{a}_k = \mathbf{0}$. Then,

$$\Pr\left[Y_{72}^{\underline{a}_k=\mathbf{0}} = 1 | \overline{\mathbf{X}}_k, Z_k = 1\right] \tag{5}$$

expresses the conditional risk of adverse neonatal or maternal outcomes 72 h after start of labor under vaginal mode of delivery only given the predictor history up to time k and that the woman is still in labor and no adverse outcome has occurred at time k (Figure 2). We retrieve Estimand 2 as the special case where the moment of intended use is restricted to $k = 0$.

Note that in this example, as well as in the next two examples, we assume a simplified situation in which the prediction model is revisited every hour to update the risks. In practice, a prediction model is revisited when a decision needs to be made, which can be at any time point since the last decision.

3.4.3 | Estimand 6: Vaginal Mode of Delivery at Moment of Prediction and Thereafter Intervention Policy "As Usual"

A consideration in sequential prediction is to define the duration of fixing the intervention strategy under which predictions are made. As explained in the discussion of Estimand (3), fixing the intervention strategy for the entire prediction window has benefits in terms of generalizability, but it might not be an intervention option that corresponds closely to actual clinical practice. In the current example, it might be preferable to fix the intervention for 1 h because the prediction model is revisited every hour.

Define the intervention option as setting mode of delivery to vaginal delivery for 1 h at time k, that is, $a_k = 0$, and follow usual intervention policy during all subsequent times, that is, $\underline{A}_{k+1}$. Then,

$$\Pr\left[Y_{72}^{a_k=0,\underline{A}_{k+1}} = 1 | \overline{\mathbf{X}}_k, Z_k = 1\right] \tag{6}$$

expresses the conditional risk of adverse neonatal or maternal outcomes 72 h after start of labor under vaginal delivery for the next hour and having the option of cesarean section at subsequent time points according to usual intervention policy. This risk is conditional on the predictor history up to time k and that the woman is still in labor and no adverse outcome has occurred at time k. The duration over which the intervention option is fixed is in this case informed by the time till next decision moment. We retrieve Estimand 3 as the special case where the moment of intended use if restricted to $k = 0$.

[2] We realize that maternal age, parity, and history of preterm birth are time-fixed predictors, but do not explicate this in the notation for simplicity.

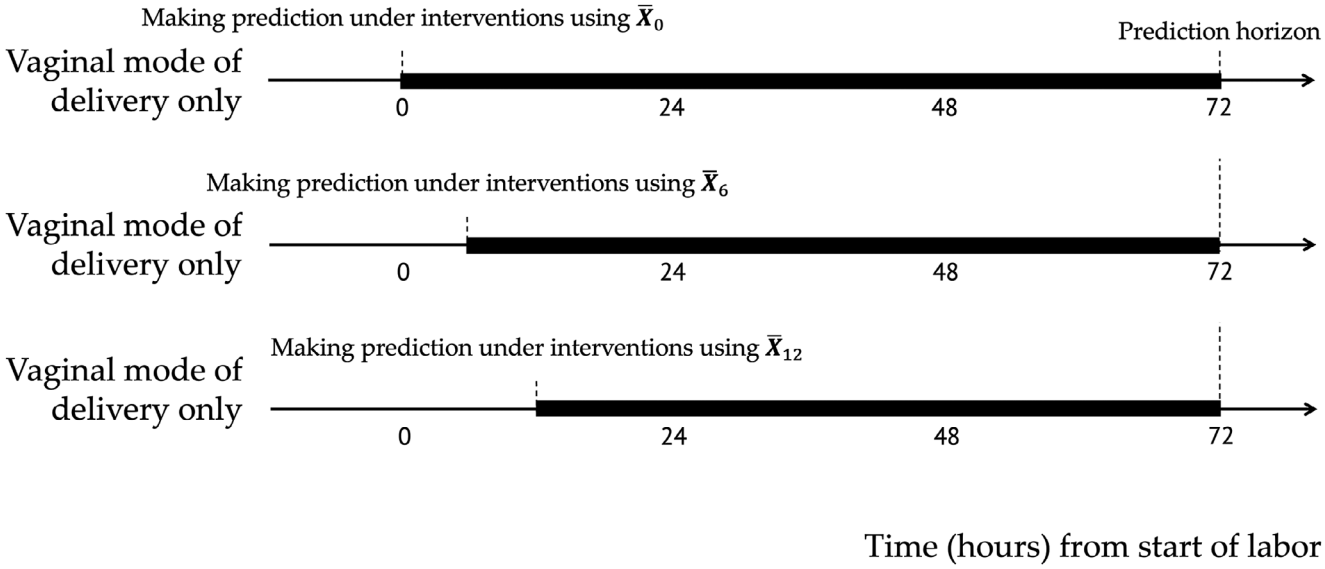

FIGURE 2 | Depiction of sequential prediction under interventions for Estimand 5. The black rectangles indicate the intervention option under which predictions are made: vaginal delivery, which is sustained throughout labor. The moment of making the prediction and prediction horizon shift over time to provide sequential updated predictions.

3.4.4 | Estimand 7: Vaginal Mode of Delivery at Moment of Prediction and a Short Prediction Horizon

Finally, estimands for sequential prediction require that the time axis of the outcome, that is, the prediction horizon, is defined. This element of the estimand determines what information the prediction model provides to inform the intervention decision. In the example, the prediction model is used to estimate risks of relatively acute outcomes (within 72 h) for the mother and child. In other cases, a long prediction horizon may be apt, for instance, when predicting (side) effects that might develop months or years later. When the predictions are revisited regularly, it might be preferable to estimate a conditional outcome risk until the next moment the prediction model is consulted. We could, for example, define the prediction horizon at $k+1$. Then,

$$\Pr\left[Y_{k+1}^{a_k=0} = 1 | \overline{\mathbf{X}}_k, Z_k = 1\right] \tag{7}$$

expresses the conditional risk of adverse neonatal or maternal outcomes within 1 h after consulting the prediction model under vaginal delivery for the next hour given the predictor history up to time k and that the woman is still in labor and no adverse outcome has occurred at time k. The duration of fixing the intervention and the prediction horizon are in this case informed by the time till next decision moment. Because of the short prediction horizon, it might also be informative to define an intermediate outcome instead, such as fetal hypoxia.

4 | Remarks on Identification and Estimation

After estimands for prediction under interventions are defined in a study, the subsequent steps are assessing identifiability given the observed data and choosing an appropriate estimator. Identifiability and estimation are beyond the scope of the current work, but we provide a few pointers here.

Assessing whether the observed information is sufficient to identify the desired estimand for prediction under interventions involves evaluating whether causal identification assumptions are met. The identifiability assumptions typically include consistency, (conditional and sequential) exchangeability, and (sequential) positivity (Boyer, Dahabreh, and Steingrimsson 2023; Hernán and Robins 2020; Keogh and Van Geloven 2024). We refer the reader to other material for discussions on how to assess identifiability, such as Boyer, Dahabreh, and Steingrimsson (2023), Hernán and Robins (2020), and Keogh and Van Geloven (2024), but make two remarks relevant to prediction under interventions in particular. The predictor(s) that are defined as part of an estimand are typically chosen based on (bedside) availability and prognostic value. They are not selected with causal assumptions in mind. Evaluation of (sequential) exchangeability might reveal that additional (time-varying) covariates need to be taken into account during estimation for confounding adjustment. This may be because the covariates available during deployment are only a subset of the adjustment set needed for confounding adjustment, also referred to as "runtime confounding" (Coston, Kennedy, and Chouldechova 2020). In addition, when targeting predictions under sustained treatment strategies, time-varying covariates are often needed to account for time-varying confounding.

Regarding (sequential) positivity, it needs to be possible to observe individuals under all formulated intervention options for all covariate patterns given by both the adjustment variables needed for confounding control and the predictors in the model (Boyer, Dahabreh, and Steingrimsson 2023). This is more stringent compared to (sequential) positivity for average treatment effects, where this requirement only conditions on the covariates needed for confounding adjustment. As a result, for predictions under interventions often more data are needed or more parametric assumptions may need to be made.

If an estimand is identifiable given observed data, we can then select an estimator from many available existing methods to

estimate the formulated estimands, mainly stemming from the causal inference literature on dynamic/conditional treatment effect estimation (Chakraborty and Moodie 2013; Gran et al. 2010; Morzywołek, Decruyenaere, and Vansteelandt 2023). Such methods can be used to develop a prediction model under the hypothetical intervention scenario(s) as if all individuals had followed the defined intervention option (Dickerman et al. 2022; Lin et al. 2021; Sperrin et al. 2018; van Geloven et al. 2020). In the sequential prediction context, additional care should be given to the way the history of the predictors up to the moment of prediction is summarized (Putter and van Houwelingen 2022; Rizopoulos 2012; Rizopoulos et al. 2023; van Houwelingen and Putter 2011). Further methodological work is needed to develop dedicated methods for estimating and validation of prediction under interventions (Keogh and Van Geloven 2024), for example, combining the selection of predictors from a set of candidate predictors while remaining sufficient adjustment for confounding. Standard estimation and validation methods do not apply when using observational data, because prediction under interventions involves obtaining predictions of the outcome under interventions that are different from those received by a subset of individuals.

5 | Discussion

The current work provides recommendations about how to formulate estimands for sequential prediction under interventions. We combined guidance for formulating estimands in prediction, clinical trial, and causal inference literature and discussed key considerations (Table 1). The considerations were illustrated by formalizing seven estimands in a clinical case study about the decision on mode of delivery in women giving birth. These insights on estimands extend previous work on predictions under interventions for single-stage decisions and link principles of dynamic prediction modeling to prediction under intervention options.

In empirical studies focusing on sequential prediction under interventions, it is likely that multiple estimands are relevant for the prediction problem at hand. How to align the moments of intended use, intervention duration, and prediction horizon should depend on context-specific knowledge and preferences of different model users. Ultimately, the selected (set of) estimand(s) should capture the risks under intervention options that provide relevant information for decision making. The choice of moments of intended use of a prediction model under interventions can also be determined by estimating the optimal timing of risk assessments (Gasperoni et al. 2023).

We discussed estimands for an example of an assistive prediction model that can provide information on risks for mother and child under the defined intervention options. These risks are one piece of information that could inform the decision on mode of delivery and leaves weighting of other sources of information to arrive at a decision to the model users. As such, *assistive* prediction models fit the process of shared decision making more naturally compared to optimal dynamic intervention rules (which are sometimes referred to as *directive* decision support tools; Reilly and Evans 2006).

The transition from outcome risk estimation to a particular intervention or triage decision is typically described to take place only after model development and validation, in a so-called impact analysis study (Moons, Altman, et al. 2009; Reilly and Evans 2006). In an impact analysis, the impact of the use of a prediction model on patient outcomes or (cost) efficiency of care is evaluated in a setting in which some health care professionals are assigned to use the prediction model when making decisions about an intervention—preferably by randomization. The current work links predictions to decisions already in the model development phase, which ensures that predictions are aligned to their intended clinical impact. We recommend that development of such prediction models always starts by defining relevant estimands. The need for other good practices in prediction modeling, like external validation and impact studies, remains (Boyer, Dahabreh, and Steingrimsson 2023; de Hond et al. 2022; Keogh and Van Geloven 2024).

Causal research closely related to prediction under interventions focuses on estimation of individualized intervention *effects* conditional on covariates with the goal to investigate heterogeneity in intervention effects in randomized controlled trials (Cai, van Buuren, and Vink 2022; Hoogland et al. 2021; Kent et al. 2020a, 2020b; Kent, Steyerberg, and van Klaveren 2018). Still, most work in this area seems to assume that an intervention decision is made at a single point in time, at the moment of randomization, which often does not align with decision making in clinical practice. Furthermore, there is debate whether a conditional difference in means should be used to assign interventions based on who benefits most (de Vries, Groenwold, and Luedtke 2020; VanderWeele et al. 2019). Our work focuses on individualized outcome risks under certain intervention options rather than estimating intervention contrasts so that the prediction model can be used as an assistive tool in making an intervention decision.

Our clinical case study did not address all intricacies of decisions to be made when defining estimands for sequential prediction under interventions. For example, others have proposed to set the intervention for the first time point and afterward assume an optimal intervention regime (Vansteelandt and Joffe 2014). The use case contained some specific features that do not occur in general such as irreversibility of the intervention "cesarean section." We specified a composite outcome, but knowing the risk of each outcome separately might be more helpful in practice to weigh the risks for the mother and neonate. The prediction horizon was set at a time point of 72 h, which was relevant in the clinical context, but may complicate the estimation when a single model is used for prediction windows of different duration. Because our focus was on defining estimands, we kept the discussion of estimation strategies to a minimum. Yet, estimating conditional outcome risks under interventions might be particularly challenging in a sequential setting because of time-varying confounding, the more stringent positivity conditions, and the need to summarize the predictor history up to the moment of intended use.

We selected one clinical setting to introduce estimands for sequential prediction under interventions, but there are many other examples of clinical settings in which sequential prediction under interventions is relevant. One possible field is organ transplantation, to inform decisions about organ allocation for individuals in need of a transplant. Every time an organ becomes

available, the risk of survival under receiving or not receiving the organ can be predicted for individuals on the wait list (Keogh and Van Geloven 2024). Another possible clinical setting is the intensive care unit, where many parameters are measured regularly to monitor patients. For example, the decision whether or not to opt for dialysis in patients with acute kidney disease can be informed by predicting survival under starting dialysis now or refraining from/delaying start of dialysis (Morzywołek 2023). Finally, when monitoring patients with a chronic disease, like diabetes, sequential prediction under interventions can be used during regular physician visits, for instance, every 6 months. Predicting the risk of adverse outcomes under continuation of the current medication or change in medication can inform prescribing decisions.

Defining a suitable estimand for sequential prediction under interventions is far from trivial, but is a pivotal starting point for development of any prediction model intended to inform intervention decisions at multiple time points. Different estimands can produce risk estimates that are relatively similar, but in some situations estimates can differ substantially (Prosepe et al. 2022; van Geloven et al. 2020). Formulating estimands prevents intervention decisions from being misguided by information from prediction models. The current work illustrates how estimands for sequential prediction under interventions can be formulated.

Conflicts of Interest

The authors declare no conflicts of interest.

Data Availability Statement

The authors state that no additional data were used for this study.

References

Boyer, C. B., I. J. Dahabreh, and J. A. Steingrimsson. 2023. "Assessing Model Performance for Counterfactual Predictions." Preprint https://doi.org/10.48550/arXiv.2308.13026.

Cai, M., S. van Buuren, and G. Vink. 2022. "How to Relate Potential Outcomes: Estimating Individual Treatment Effects Under a Given Specified Partial Correlation." Preprint https://doi.org/10.48550/arXiv.2208.12931.

Chakraborty, B., and E. E. Moodie. 2013. *Statistical Methods for Dynamic Treatment Regimes: Reinforcement Learning, Causal Inference, and Personalized Medicine*. Vol. 76. New York: Springer.

Coston, A., E. H. Kennedy, and A. Chouldechova. 2020. "Counterfactual Predictions Under Runtime Confounding." *Advances in Neural Information Processing Systems* 33: 4150–4162.

Coston, A., A. Mishler, E. H. Kennedy, and A. Chouldechova. 2020. "Counterfactual Risk Assessments, Evaluation, and Fairness." In *Proceedings of the 2020 Conference on Fairness, Accountability, and Transparency*, 582–593.

Dawid, A. P., and V. Didelez. 2010. "Identifying the Consequences of Dynamic Treatment Strategies: A Decision Theoretic Overview." *Statistics Surveys* 4: 184–231.

de Hond, A. A., A. M. Leeuwenberg, L. Hooft, et al. 2022. "Guidelines and Quality Criteria for Artificial Intelligence-Based Prediction Models in Healthcare: A Scoping Review." *NPJ Digital Medicine* 5, no. 1: 2.

de Vries, B. B. L. P., R. H. H. Groenwold, and A. Luedtke. 2020. "Re. Selecting Optimal Subgroups for Treatment Using Many Covariates." *Epidemiology* 31, no. 4: e33–e34.

Dickerman, B. A., I. J. Dahabreh, K. V. Cantos, et al. 2022. "Predicting Counterfactual Risks Under Hypothetical Treatment Strategies: An Application to HIV." *European Journal of Epidemiology* 37, no. 4: 367–376.

Dickerman, B. A., and M. A. Hernán. 2020. "Counterfactual Prediction is Not Only for Causal Inference." *European Journal of Epidemiology* 35: 615–617.

European Medicines Agency. 2020. "ICH E9 (R1) Addendum on Estimands and Sensitivity Analysis in Clinical Trials to the Guideline on Statistical Principles for Clinical Trials." Accessed September 12, 2023. https://www.ema.europa.eu/en/documents/scientific-guideline/ich-e9-r1-addendum-estimands-sensitivity-analysis-clinical-trials-guideline-statistical-principles_en.pdf.

Gasperoni, F., C. H. Jackson, A. M. Wood, et al. 2023. "Optimal Risk-Assessment Scheduling for Primary Prevention of Cardiovascular Disease." Preprint https://doi.org/10.48550/arXiv.2302.04992.

Goetghebeur, E., S. le Cessie, B. de Stavola, E. E. Moodie, I. Waernbaum, and The Topic Group Causal Inference (TG7) of the STRATOS Initiative. 2020. "Formulating Causal Questions and Principled Statistical Answers." *Statistics in Medicine* 39, no. 30: 4922–4948.

Gran, J. M., K. Røysland, M. Wolbers, et al. 2010. "A Sequential Cox Approach for Estimating the Causal Effect of Treatment in the Presence of Time-Dependent Confounding Applied to Data From the Swiss HIV Cohort Study." *Statistics in Medicine* 29, no. 26: 2757–2768.

Hernán, M. A., J. Hsu, and B. Healy. 2019. "A Second Chance to Get Causal Inference Right: A Classification of Data Science Tasks." *Chance* 32, no. 1: 42–49.

Hernán, M. A., and J. M. Robins. 2020. *Causal Inference: What If*. Boca Raton: Chapman & Hall/CRC.

Hoogland, J., J. IntHout, M. Belias, et al. 2021. "A Tutorial on Individualized Treatment Effect Prediction From Randomized Trials With a Binary Endpoint." *Statistics in Medicine* 40, no. 26: 5961–5981.

Kent, D. M., J. K. Paulus, D. van Klaveren, et al. 2020. "The Predictive Approaches to Treatment Effect Heterogeneity (PATH) Statement." *Annals of Internal Medicine* 172, no. 1: 35–45.

Kent, D. M., E. Steyerberg, and D. van Klaveren. 2018. "Personalized Evidence Based Medicine: Predictive Approaches to Heterogeneous Treatment Effects." *BMJ* 363–k4245.

Kent, D. M., D. van Klaveren, J. K. Paulus, et al. 2020. "The Predictive Approaches to Treatment Effect Heterogeneity (PATH) Statement: Explanation and Elaboration." *Annals of Internal Medicine* 172, no. 1: W1–W25.

Keogh, R. H., and N. Van Geloven. 2024. "Prediction Under Interventions: Evaluation of Counterfactual Performance Using Longitudinal Observational Data." *Epidemiology* 35, no. 3: 329–339.

Knaus, M. C., M. Lechner, and A. Strittmatter. 2021. "Machine Learning Estimation of Heterogeneous Causal Effects: Empirical Monte Carlo Evidence." *Econometrics Journal* 24, no. 1: 134–161.

Lin, L., M. Sperrin, D. A. Jenkins, G. P. Martin, and N. Peek. 2021. "A Scoping Review of Causal Methods Enabling Predictions Under Hypothetical Interventions." *Diagnostic and Prognostic Research* 5: 1–16.

Luijken, K., R. H. Groenwold, B. Van Calster, E. W. Steyerberg, and M. van Smeden. 2019. "Impact of Predictor Measurement Heterogeneity Across Settings on the Performance of Prediction Models: A Measurement Error Perspective." *Statistics in Medicine* 38, no. 18: 3444–3459.

Moodie, E. E., T. S. Richardson, and D. A. Stephens. 2007. "Demystifying Optimal Dynamic Treatment Regimes." *Biometrics* 63, no. 2: 447–455.

Moons, K. G., D. G. Altman, J. B. Reitsma, et al. 2015. "Transparent Reporting of a Multivariable Prediction Model for Individual Prognosis

or Diagnosis (TRIPOD): Explanation and Elaboration." *Annals of Internal Medicine* 162, no. 1: W1–W73.

Moons, K. G., D. G. Altman, Y. Vergouwe, and P. Royston. 2009. "Prognosis and Prognostic Research: Application and Impact of Prognostic Models in Clinical Practice." *BMJ* 338: b604.

Moons, K. G., J. A. de Groot, W. Bouwmeester, et al. 2014. "Critical Appraisal and Data Extraction for Systematic Reviews of Prediction Modelling Studies: The CHARMS Checklist." *PLoS Medicine* 11, no. 10: e1001744.

Moons, K. G., P. Royston, Y. Vergouwe, D. E. Grobbee, and D. G. Altman. 2009. "Prognosis and Prognostic Research: What, Why, and How?" *BMJ* 338: b375.

Morzywołek, P. 2023. "Causal Inference Methods to Optimize Clinical Decision-Making in Treatment Initiation Based on Routinely Collected Data." PhD thesis, Faculty of Sciences, Ghent University.

Morzywołek, P., J. Decruyenaere, and S. Vansteelandt. 2023. "On Weighted Orthogonal Learners for Heterogeneous Treatment Effects." Preprint https://doi.org/10.48550/arXiv.2303.12687.

Murphy, S. A. 2003. "Optimal Dynamic Treatment Regimes." *Journal of the Royal Statistical Society Series B: Statistical Methodology* 65, no. 2: 331–355.

National Institute for Health and Care Excellence. 2019. "Intrapartum Care for Women With Existing Medical Conditions or Obstetric Complications and Their Babies." https://www.nice.org.uk/guidance/ng121.

National Institute for Health and Care Excellence. 2021. "Caesarean Birth." https://www.nice.org.uk/guidance/ng192.

National Institute for Health and Care Excellence. 2023. "Intrapartum Care." https://www.nice.org.uk/guidance/ng235.

Petersen, M. L., and M. J. van der Laan. 2014. "Causal Models and Learning From Data: Integrating Causal Modeling and Statistical Estimation." *Epidemiology* 25, no. 3: 418–426.

Prosepe, I., R. H. Groenwold, R. Knevel, R. Pajouheshnia, and N. van Geloven. 2022. "The Disconnect Between Development and Intended Use of Clinical Prediction Models for Covid-19: A Systematic Review and Real-World Data Illustration." *Frontiers in Epidemiology* 2: 899589.

Putter, H., and H. C. van Houwelingen. 2022. "Landmarking 2.0: Bridging the Gap Between Joint Models and Landmarking." *Statistics in Medicine* 41, no. 11: 1901–1917.

Reilly, B. M., and A. T. Evans. 2006. "Translating Clinical Research into Clinical Practice: Impact of Using Prediction Rules to Make Decisions." *Annals of Internal Medicine* 144, no. 3: 201–209.

Riley, R. D., D. van der Windt, P. Croft, and K. G. Moons. 2019. *Prognosis Research in Healthcare: Concepts, Methods, and Impact*. United Kingdom: Oxford University Press.

Rizopoulos, D. 2012. *Joint Models for Longitudinal and Time-to-Event Data: With Applications in R*. United States of America: CRC Press.

Rizopoulos, D., J. M. Taylor, G. Papageorgiou, and T. M. Morgan. 2023. "Using Joint Models for Longitudinal and Time-to-Event Data to Investigate the Causal Effect of Salvage Therapy after Prostatectomy." Preprint https://doi.org/10.48550/arXiv.2309.02115.

Robins, J. M. 2004. "Optimal Structural Nested Models for Optimal Sequential Decisions." In *Proceedings of the Second Seattle Symposium in Biostatistics: Analysis of Correlated Data*, 189–326. New York: Springer.

Schuit, E., A. Kwee, M. E. Westerhuis, et al. 2012. "A Clinical Prediction Model to Assess the Risk of Operative Delivery." *BJOG: An International Journal of Obstetrics & Gynaecology* 119, no. 8: 915–923.

Sperrin, M., D. Jenkins, G. P. Martin, and N. Peek. 2019. "Explicit Causal Reasoning is Needed to Prevent Prognostic Models Being Victims of Their Own Success." *Journal of the American Medical Informatics Association* 26, no. 12: 1675–1676.

Sperrin, M., G. P. Martin, A. Pate, T. Van Staa, N. Peek, and I. Buchan. 2018. "Using Marginal Structural Models to Adjust for Treatment Drop-in When Developing Clinical Prediction Models." *Statistics in Medicine* 37, no. 28: 4142–4154.

Sterne, J. A., M. A. Hernán, B. C. Reeves, et al. 2016. "ROBINS-I: A Tool for Assessing Risk of Bias in Non-Randomised Studies of Interventions." *BMJ* 355: i4919.

Steyerberg, E. W. 2009. *Clinical Prediction Models: A Practical Approach to Development, Validation, and Updating*. New York: Springer.

van Geloven, N., R. H. Keogh, W. van Amsterdam, et al. 2024. Causal blind spots when using prediction models for treatment decisions. arXiv preprint arXiv:2402.17366.

van Geloven, N., S. A. Swanson, C. L. Ramspek, et al. 2020. "Prediction Meets Causal Inference: The Role of Treatment in Clinical Prediction Models." *European Journal of Epidemiology* 35: 619–630.

van Houwelingen, H., and H. Putter. 2011. *Dynamic Prediction in Clinical Survival Analysis*. United States of America: CRC Press.

VanderWeele, T. J., A. R. Luedtke, M. J. van der Laan, and R. C. Kessler. 2019. "Selecting Optimal Subgroups for Treatment Using Many Covariates." *Epidemiology* 30, no. 3: 334–341.

Vansteelandt, S., and M. Joffe. 2014. "Structural Nested Models and G-Estimation: The Partially Realized Promise." *Statistical Science* 29, no. 4: 707–731.

Whittle, R., K.-L. Royle, K. P. Jordan, R. D. Riley, C. D. Mallen, and G. Peat. 2017. "Prognosis Research Ideally Should Measure Time-Varying Predictors at Their Intended Moment of Use." *Diagnostic and Prognostic Research* 1: 1–9.