8

Assessing dynamic treatment strategies

Carlo Berzuini, A. Philip Dawid
Statistical Laboratory, Centre for Mathematical Sciences, University of Cambridge, Cambridge, UK
Vanessa Didelez
Department of Mathematics, University of Bristol, Bristol, UK

8.1 Introduction

We continue the discussion of sequential data-gathering and decision-making processes, started in the preceding chapter in this volume. The archetypical context is that of a sequence of medical decisions, taken at different time points during the follow-up of the patient, each decision involving choice of a treatment in the light of any interim responses or adverse reactions to earlier treatments. The problem consists of predicting the consequences that a (possibly new and possibly hypothetical) treatment plan will have on a future patient, by learning from the performances of past medical decision makers on past patients. While we make constant reference to a medical application context, the scope of the method is of much broader relevance.

Our approach to this problem owes immensely to the seminal work of James Robins. In a rich series of papers (1986, 1987, 1988, 1989, 1992, 1997, 2000, 2004), Robins and Wasserman (1997) and Robins *et al*. (1999), Robins introduces the idea of different hypothetical studies (these studies being analogous to our notion of 'regimes'), in which different treatment strategies are applied. He examines conditions on the relationships between these studies, under which those treatment strategies can be compared. Among these, the *sequential randomization* condition is closely related to the 'stability' assumption used in this chapter. Furthermore, Robins (1986) introduced the *G*-computation algorithm (a special version of

Causality: Statistical Perspectives and Applications, First Edition. Edited by Carlo Berzuini, Philip Dawid and Luisa Bernardinelli.

dynamic programming) to evaluate a sequential treatment strategy, which works where standard regression – even under the sequential randomization assumption – fails. Most of the concepts expounded in this chapter have a counterpart in Robin's work, although, in the interest of readability, we shall frequently abstain from making the relationships with his work explicit.

We shall assume that we have an idea of which variables informed past decisions, without assuming that these variables have all been observed in the past. Nor shall we assume that the rules that governed past decisions are discernible or similar to those that will guide the future plan of interest.

In contrast with the previous chapter, we do use causal diagrams. In the context of sequential plan identification, causal diagrams (or, better, a particular form of causal diagram) were first advocated by Pearl and Robins (1995). In accord with these authors, we assume that the relevant causal knowledge is encoded in the form of a diagram with a completely specified topology, whose associated numerical conditional probabilities are given only for a *subset* of its variables, so-called *observed* variables. The remaining variables in the diagram are 'unobserved' or 'latent'. We shall, however, introduce a form of graphical representation of causality that differs from Pearl and Robins's diagrams in some respects. Part of the problem will be to characterize situations in which the estimation of the causal effects of the treatment plan of interest is not invalidated by unobserved confounding introduced by the latent variables. Throughout this chapter we use previous results by Dawid and Didelez (2010).

Elja Arjas, both in his chapter (Chapter 7) in this volume and in his work with Parner (2004), considers two probability models. One (called the 'obs' model) models the observational study that has generated the data. The other (called the 'ex' model) models the consequences of the future (perhaps hypothetical) application of the treatment plan of interest. In our approach, this distinction is embodied in a sort of decision parameter, called the 'regime indicator', which indicates the particular regime, 'obs' or 'ex', in operation. Supplementing the set of domain variables with the mentioned regime indicator results in an 'augmented' set of variables. In our approach, conditions for the equivalence of Arjas's 'obs' and 'ex' distributions are phrased in terms of conditional independence conditions on this augmented set. Both chapters, ours and that of Arjas, avoid formulations based on ideas of potential outcomes or counterfactuals. We illustrate the methods and their motivations with the aid of two examples. Following Robins, we choose the first example in the area of the treatment of HIV infection. The second example will be in the treatment of abdominal aortic aneurysm.

This chapter is more about the identifiability of treatment plan effects, than about the methods of estimation and computation one is supposed to use if the problem turns out to be identifiable. The latter problems are discussed in greater detail in Chapter 17 in this volume, by Daniel, De Stavola and Cousens, who illustrate their method with the aid of the same HIV example we use here. An introduction to this HIV example is given in the next section.

8.2 Motivating example

We start with an example in human immunodeficiency virus (HIV) infection. The standard treatment of this disease involves an *initiation therapy* based on a combination of three antiretroviral agents, chosen from a larger pool of candidate drugs. In many patients, this initial therapy will eventually *fail*. Failure criteria include virologic evidence, for example a decrease by less than a factor of 10 in the virus RNA level during the last eight weeks or

development of life-threatening toxic effects or, for example, dislipidaemia. Therapy failure will prompt consideration of a possible replacement of one or more components of the initial drug cocktail. The decision whether to do the replacement, of the component to replace, and of the replacement component, will be based on any available information about the past evolution of the disease and of the virus (possibly also on the basis of viral genome sequencing information). Various types of decision are involved, such as: what threshold decrease in viral RNA should be interpreted as a therapy failure, and/or to prompt therapy substitution, and, in the latter case, what criterion should we use to choose the replacement drug, how many components of the drug cocktail should be chosen and which ones etc.? The rules that govern these decisions (whether originating from the doctor's experience or medical guidelines) constitute the *treatment strategy*.

One possible strategy consists of applying a fixed, pre-defined, sequence of drug cocktails and dosages, where the switches from one cocktail to the next occur at pre-established times, irrespective of any incoming information about the patient. This is called a *static* strategy, or plan. More interesting and relevant are *dynamic* treatment plans, where the treatment decisions adapt to, and depend on, the accruing series of patient events. Treatment plans are usually compared in terms of their impact on the patient's disease course, for example in terms of the length of the AIDS-free survival time, although cost-effectiveness considerations may easily be brought into the framework (we shall not say much about this). This chapter addresses the problem of predicting the consequences on a hypothetical future patient of a treatment plan of interest, without requiring this plan to have ever been considered or implemented in the past. Application of the method to alternative treatment plans provides a basis for *comparing* them, as well as a basis for determining the *best* plan under given constraints, the so-called *optimization* problem. A discussion of methodological and computational issues in optimization is outside the scope of the present chapter.

8.3 Descriptive versus causal inference

Standard statistical software offers tools to describe the endpoint experience of groups of patients subjected to different treatment strategies. Application of these tools to the data might, for example, reveal that 'patients that have been treated according to strategy 1 have fared better than those that have been treated according to strategy 2'.

The problem with this 'descriptive' approach is that – while possibly useful and correct in a sense – it may be misleading both scientifically and in terms of practical (medical) implications. What the clinicians and the scientists want is to predict the effect of assigning (by intervention) a specific treatment to a *future* case, at any time of his/her evolution, and conditional on the gradually accumulating information we collect from him/her. Can this knowledge be extracted from observational data, collected from past medical records?

The problem is that future cases may differ from those in the data in various important ways, for example because, in the observed data, the doctor's choices of actions made use of additional – unrecorded – private information. In Arjas's terminology, this would correspond to the causal variables, $\{A_k\}$, *not* being unconfounded relative to the latent variables, in the sense of his Definition 2. Moving in a similar direction as Arjas, in this chapter we formalize the requirements under which past empirical evidence can be validly carried into future decisions. Whenever these requirements appear unreasonable in a real application, we must be prepared to conclude that the available data have little practical implication with respect to future (medical)

practice. While often useful, descriptive analysis may be misleading if used to corroborate causal claims, in those situations where the formal requirements discussed in this chapter are not carefully negotiated.

Of course, if we were able to carry out a large experiment in which sample patients are randomized over the set of treatment strategies of interest, then difficulties would, at least in part, disappear. Confounding would be eliminated by randomization. However, in very many circumstances, we shall have to (or wish to) exploit data gathered under purely observational circumstances, and this motivates our effort.

8.4 Notation and problem definition

Consider the ordered sequence of time points $\{T_k\}$ (for $k = 0, 1, \ldots, N+1$), with $0 \equiv T_0 < T_1 < T_2 < \cdots < T_{N+1}$. Unlike in the preceding chapter, we take this sequence to be fixed. Moreover, for simplicity, we shall take it to be the same in all patients. In a medical treatment decision context, these time points will usually be defined on a patient-specific time scale, such as, for example, time since diagnosis, rather than, say, calendar year. For $k = 1, \ldots, N+1$, we consider the following patient-specific variables:

X_k: the values of a set of covariates observed at T_k,

A_k: an action or treatment decision, performed right after T_k (undefined for $k = N+1$). In certain applications, this variable could denote the level of a certain exposure at time T_k.

Our assumption that the domain variables are completely ordered in time is not essential, but it will make our discussion simpler. Writing, for example (A_1, A_2) for $A_1 \cup A_2$, we let the symbol $\overline{A}_k$ denote the partial sequence $(A_o, A_1, \ldots, A_k)$ of actions performed up to time T_k, and a similar notation be used for other variables in the problem. The $\{X_k\}$ sequence contains a distinguished variable, $Y \equiv X_{N+1}$, called the *endpoint* variable. This could be the state of health of the patient at the end of a given time period or, for example, a summary of the entire realized history, $(\overline{X}_N, \overline{A}_N)$, of the process. This will cover situations in which alternative plans are compared in terms of the involved treatment costs.

Throughout this chapter, we are assuming our data to have been collected under an *observational* regime, which we denote as o. Such a regime will generally involve the passive observation of the performance of past decision makers. One example is in situations where the data are retrieved from medical records describing certain treatments applied to certain patients, together with the disease courses of these patients. In our notation, these data constitute sample realizations (or 'paths') of the $(\overline{X}_{N+1}, \overline{A}_N)$ process. The method we are going to describe does not require us to discern the decision criteria in operation under the observational regime, nor the source of these criteria, be it the doctor's personal experience, or public guidelines, imposed protocols, etc. In contrast with the observational regime, we have the *experimental* regime, denoted by e, in which a hypothetical future patient is treated according to a specific (static or dynamic) plan, which we wish to evaluate. Such a plan will consist of a specific set of rules for determining the value of each A_k, at time T_k, possibly on the basis of information about $(\overline{A}_{k-1}, \overline{X}_{k-1})$. Part of the notion of the 'experimental regime' also the modalities with which each treatment action will be carried out and, more in general, the environmental and technological conditions and the medical expertise, which will surround the application of these actions. The aim of the method is to predict the effect of the experimental

regime e on Y, on the basis of the data collected under o. We do not need to assume that the decision criteria that inform the plan under assessment have ever been adopted in the past.

We introduce a nonrandom *regime indicator*, σ, taking values in (o, e). We think of the value of σ as being determined externally, before any observations on $(\overline{X}_{N+1}, \overline{A}_N)$ are made. All probability statements about $(\overline{X}_{N+1}, \overline{A}_N)$ must then be explicitly or implicitly conditional on the value of σ. We use the symbol $p(Y \mid \sigma = e)$ to denote the marginal density of Y for a future patient under treatment regime e. If we are able to estimate this density, we shall be able to say something about how effective (in terms of whatever loss function we adopt) is the treatment plan under investigation. Of interest are then the conditions under which the past data, collected under regime o, can be used to determine $p(Y \mid \sigma = e)$. These we shall refer to as *plan identifiability conditions*. For simplicity, in the following, we suppose that the time grid, $\{T_k\}$, and the $(\overline{X}_{N+1}, \overline{A}_N)$ variables are the same for both regimes, o and e. For the time being, we shall assume that treatment decisions under e are influenced by the same variables that have influenced the corresponding decisions under o.

8.5 HIV example continued

With reference to our HIV example, the sequence $T_k, k = 1, \ldots, N$, might represent a fixed set of times at which the patient is visited, with T_0 representing some clinically meaningful temporal origin. The X_k covariate might represent the observed values, at T_k, of viral RNA levels, CD4 levels, AIDS status and – more broadly – safety-related variables or indicators of possible unfavourable side effects of the therapy. Immediately after information X_k has been collected, a therapeutic action, A_k, will be performed on the patient. This might consist of switching from the current drug cocktail to a new one, by a rational replacement of one or more components, in response to the past clinical evolution pattern of the patient.

8.6 Latent variables

In most situations, we shall need to introduce in the problem additional variables that are essential to describe the causal structure of the problem, but are unobserved. These variables we call 'unobserved', or *latent*, and collectively denote with the symbol $\mathcal{L}$. We shall let elements of this set be denoted by the generic symbols U and W. We shall have these variables time-subscripted, with U_k, say, characterizing the set of a relevant latent process right before time T_k and after time T_{k-1}. The symbol $\overline{L}_i$ will denote the set of all $\mathcal{L}$ variables with a subscript equal or lower than i.

In the HIV example, we may use latent variables to represent, for example, patient-specific, possibly inheritable, patterns of response to drugs. In Chapter 17, for example, in their HIV example, Daniel and colleagues introduce latent variables that categorize the patients into 'types', to indicate lower-than-average responsiveness to specific therapies and/or vulnerability to the side effects of a specific therapy (see Chapter 17 for details). Our formulation allows these variables to be organized in a time process, which allows the model to acknowledge the tendency of some patients to switch type during the course of therapy. The introduction of latent variables will result in a more detailed causal structure, and by a correspondingly more complex causal diagram representation of the problem. This has advantages. For one, it will often help justify the necessary assumptions. The stability assumption of the next section, for

example, is often easier to justify in terms of local properties of a detailed causal diagram that includes latent variables, than in terms of relationships between the observed variables of direct interest.

8.7 Conditions for sequential plan identifiability

We are now going to explore conditions for sequential plan identifiability, that is conditions under which our data, which have been collected under regime o, can be used to infer the distribution $p(Y \mid \sigma = e)$ of the endpoint variable under regime e. In general, this will *not* be possible, one reason being that, in general, the probability distribution of the observed variables will differ from one regime to the other. We have identifiability if certain properties of equivalence between the distributions of the observed variables under the two regimes are satisfied. These properties are discussed in the following.

Under the experimental regime, denoted by $\sigma = e$, the joint distribution of the observed variables factorizes into the product of the conditional distributions of each of these variables given the earlier observed variables:

$$p(y, \overline{x}_N, \overline{a}_N \mid \sigma = e) = \left\{ \prod_{k=0}^{N+1} p(x_k \mid \overline{x}_{k-1}, \overline{a}_{i-1}\,; \sigma = e) \right\} \times \left\{ \prod_{k=0}^{N} p(a_k \mid \overline{x}_k, \overline{a}_{k-1}\,; \sigma = e) \right\} \tag{8.1}$$

where $X_{N+1} \equiv Y$ and where $p(x_k \mid \overline{x}_{k-1}, \overline{a}_{k-1}\,; \sigma = e)$ really stands for

$$\int_{dom(\mathcal{L}_{k-1})} p(x_k \mid \overline{x}_{k-1}, \overline{a}_{k-1}, \overline{l}_{k-1} \mid \sigma = e)\ \mathrm{d}p(\mathcal{L}_{k-1} \mid \overline{x}_{k-1}, \overline{a}_{k-1}, \sigma = e) \tag{8.2}$$

where $dom(\mathcal{L}_{k-1})$ stands for the domain of the L_{k-1} variables. Note that the second factor on the right-hand side of Equation (8.1) represents the pre-established treatment strategy we want to assess and, as a consequence, is to be considered known a priori. Note also that we are assuming each variable to have the same property, of being observed or unobserved, under both regimes. Our inferential problem is solved if we are able, on the basis of the data collected under o, to estimate the probability distribution (8.1). We shall now see that we may indeed be able to do so if the properties of stability and positivity, discussed in the following, are satisfied.

8.7.1 Stability

In general, for a generic k, the conditional distribution $p(X_k \mid \overline{X}_{k-1}, \overline{A}_{k-1})$ will vary across regimes; that is, conditional on a particular *observed* history of the individual, the effect of a given sequence of actions s will, in general, differ systematically from one regime to the other. In our HIV example, this could be due, for example, to the fact that the doctors of one regime have a better visibility of the underlying patient's 'type' than those of the other regime. As a consequence, those patients of one regime who underwent the sequence s are not comparable (despite the conditioning on past observed history) with those patients of the other regime who underwent the same sequence. In those situations where, by contrast, the distribution $p(X_k \mid \overline{X}_{k-1}, \overline{A}_{k-1})$ does *not* differ from one regime to the other, we say we have 'stability'.

More formally, we say that the *stability* condition is satisfied if, for $k = 0, \ldots, N+1$, the following identity holds:

$$p(x_k \mid \overline{x}_{k-1}, \overline{a}_{k-1}\,;\, \sigma = e) = p(x_k \mid \overline{x}_{k-1}, \overline{a}_{k-1}\,;\, \sigma = o) \tag{8.3}$$

whenever the conditioning event, $\overline{x}_{k-1}, \overline{a}_{k-1}$, has positive probability under either of regimes o and e. Our definition of stability is closely related to Robin's concept (1986) of sequential randomization (1986). Having previously defined σ to take values in (o, e), the stability condition is equivalent to the following conditional independence property:

$$X_k \perp\!\!\!\perp \sigma \mid (\overline{X}_{k-1}, \overline{A}_{k-1}) \quad (k = 0, \ldots, N+1) \tag{8.4}$$

Here and throughout, we use the notation and theory of conditional independence introduced by Dawid (1979) as generalized in Dawid (2002) to apply also to problems involving decision or parameter variables.

A major *caveat* is the fact that, because stability is a property of the relationship between different regimes, it can never be empirically established on the basis of data collected under the sole observational regime. In order to judge how realistic the stability assumption is in a specific application, we shall need genuine insight into the subject matter. In the next section, with reference to our HIV example, we shall see that a detailed causal diagram of the problem, which includes the relevant latent variables, may help in this task. One reason is that it may be relatively easy to justify assumptions about local properties of the diagram from which, ultimately, stability can be deduced via the usual conditional independence semantics of DAGs.

Intuitively, under the stability condition (8.4), or under the equivalent condition (8.3), one can formally replace the distribution $p(x_i \mid \overline{x}_{i-1}, \overline{a}_{i-1}\,;\, \sigma = e)$ in Equation (8.1) by $p(x_i \mid \overline{x}_{i-1}, \overline{a}_{i-1}\,;\, \sigma = o)$, and thus hope to be able to estimate the first factor on the right-hand side of Equation (8.1) from the available observational data. In this case, because the second factor on the right-hand side of Equation (8.1) is known a priori, one would hope to be able to estimate the distribution $p(y, \overline{x}_N, \overline{a}_N \mid \sigma = e)$ from the data. However, some extra care is needed here. We shall – in general – have to invoke a further property called *positivity*, which is discussed in the following.

8.7.2 Positivity

If, for example, we want to assess the consequence of an interventional strategy e, under which a particular action sequence $\overline{a}^*$ may arise, we will be unable to do so if, under the observational regime that has generated our data, that particular sequence of actions arises with probability zero. Define the *positivity* condition to be satisfied when, for any strategy e and any event E defined in terms of $(\overline{X}_N, \overline{A}_N, Y)$, the inequality $p(E \mid e) > 0$ implies $p(E \mid o) > 0$. In our HIV example, positivity implies that for any combined sequence of patient events and medical actions that may arise under the treatment strategy e that we wish to evaluate, there is a non-null probability of that particular sequence arising in the observational data. An extreme example of violated positivity occurs when the regime e contemplates the possible administration of a drug molecule that was not available in the observational study that generated the data. More subtle examples of violated positivity may arise.

When both stability and positivity are satisfied, we should be able to estimate the joint distribution (8.1) on the basis of the observational data. Then, on the basis of such distribution,

we should – at least in principle – be able to obtain $p(Y \mid \sigma = e)$ by marginalization with respect to all variables but $x_{N+1} \equiv Y$. Note, however, that this procedure is extremely cumbersome from a computational point of view. A computationally more efficient procedure is discussed later in this chapter.

8.8 Graphical representations of dynamic plans

As stated in the introduction section, we consider the use of causal diagrams to be a very helpful ingredient of our method. We shall here restrict ourselves to diagrams that have the form of a directed acyclic graph (DAG). The particular class of causal diagrams we advocate, called an 'augmented causal DAG' (ADAG), differs in some respects from the graphical representations of NPSEMs discussed in Chapter 3 in this volume, and from the causal diagrams of Pearl (2009). In this section, we explain the differences between ADAGs and other graphical representations of causality, and the reasons why we favour the former.

A 'causal DAG' is supposed to model causal relations. Unlike conditional independence, which is an unambiguous property of a probability distribution, causal relations lack a clear, mathematical, definition. A DAG representing causal properties should, in principle, be something totally different from a DAG representing conditional independence properties. In the former case, the arrows are supposed to have a direct interpretation in terms of cause and effect (whatever this means) whereas for conditional independence the arrows are nothing but incidental construction features supporting the d-separation semantics described by Shpitser in Chapter 3. In spite of all this, the idea of using DAGs that simultaneously represent conditional independence and causal properties (the latter understood as describing the effects of interventions), introduced by Judea Pearl and described in his book (2009), has been extremely fruitful.

Pearl's graphical representation applies to a collection of variables measured on some system, such that we can intervene (or at least can conceive of the possibility of intervening) on any one variable or collection of variables, so as to *set* the value(s) of the associated variable(s) in a way that is determined entirely externally. This gives rise to a wide variety of interventional regimes. The observational regime corresponds to the special case where no variables are set by intervention. From a probabilistic point of view, any instance of a Pearl causal DAG coherently represents a family of joint probability distributions, consisting of those distributions that satisfy the conditional independence properties one can read off the graph, for example via d-separation. From a causal point of view, the same graph represents the 'modularity' assumption that, for any node i of the graph, its conditional distribution, given its DAG parents, is the same, no matter which variables in the system (other than i itself) are intervened upon.

An important idea, previously discussed in Chapter 4, is to enrich the causal DAG in such a way that it expresses relationships between the probabilistic behaviours of the domain variables, across the regimes (Dawid, 2010). One way of doing so, which has been inspired by influence diagrams (Dawid, 2002, 2003), is to supplement the DAG with one or more nonrandom *regime indicators*, which we shall draw as squares. One example of a regime indicator, denoted by σ, has been previously introduced in this chapter to distinguish between regimes e and o. By incorporating σ in the causal diagram, we are able to express the identity of the conditional distribution of certain sets of domain variables, given certain other sets of domain variables, between the mentioned regimes.

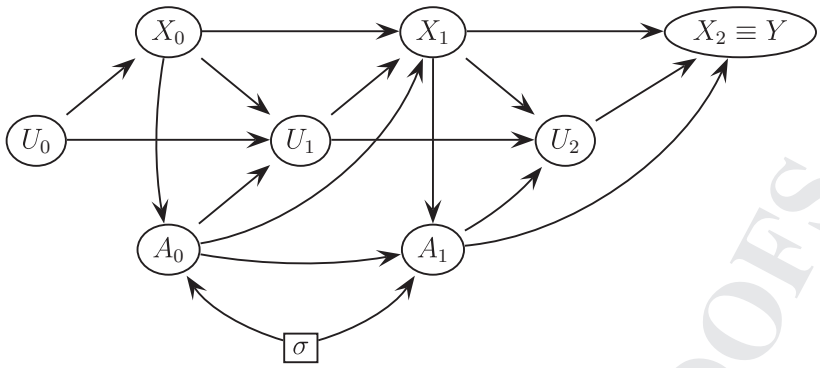

Figure 8.1 Graphical representation of the HIV problem example, with $N = 1$.

This idea is illustrated in Figure 8.1 by the ADAG representation of our HIV example, for the special case of $N = 1$ to preserve visual simplicity. The diagram uses the variable symbols introduced in Section 8.5, except for the latent process (U_0, U_1). In line with the way this example is dealt with in Chapter 17 in this volume, take U_k to characterize the patient just before time T_k as belonging to one of a number of 'types', in the sense of Section 8.6. Now, the obtained diagram satisfies the stability condition (8.4), as can be checked by applying the usual d-separation semantics of DAGs to it. In the graph of Figure 8.1, stability is a consequence of the following two topological properties:

1. $\sigma \rightarrow (X, U)$ arrows are missing. This means that, conditional on past history, the effect of any particular sequence of actions, $(A_0 = a_0, A_1 = a_1, \ldots)$, on the patient is the same, regardless of the particular regime in operation
2. $U \rightarrow A$ arrows are missing; that is, under each regime, and at any stage of the decision process, treatment assignments are determined by some deterministic, or randomizing, device, which only has the values of earlier *observed* (under both regimes) variables as inputs.

Property 1 is the graphical counterpart of what Dawid and Didelez (2010) call the *extended stability* condition. Property 2 is directly related to Robin's concept (1986) of sequential randomization. The conditional independence version of property 2, for $k = 1, \ldots, N$ and for $\sigma \in (o, e)$, is $A_k \perp\!\!\!\perp \overline{U}_{k-1} \mid \overline{X}_k, \overline{A}_{k-1}, \sigma$. This appears to be a discrete-time version of Definition 2 (ii) of Arjas and Parner (2004). We note that property 2, if not accompanied by property 1, does not per se guarantee stability. We also note that validity of the two properties above is a sufficient, but not necessary condition for stability. Further combinations of properties that result in stability are discussed in X?.

As noted by Daniel and colleagues in Chapter 17, property 1 appears to be plausible in our HIV example. The reason is that the effect of taking a particular anti-retroviral cocktail is likely be the same irrespective of whether the therapy is assigned by a doctor within the observational study o or assigned by some other doctor (assuming, of course, that no major changes in the selective resistance of viral agents have occurred between regimes o and e). Property 2 is realistic if no information about the latent patient's type was available to the

doctors under both regimes, except possibly indirectly through the observed patient's covariate process. We conclude that, under the assumptions represented by the graph of Figure 8.1, and if the positivity condition is also valid, the data collected from medical records of past HIV patients can be used to predict the relative effectiveness of alternative treatment strategies on future patients. In the next section we introduce a second example of study, in the field of abdominal aortic aneurysm.

8.9 Abdominal aortic aneurysm surveillance

An abdominal aortic aneurysm (AAA) is a serious vascular disease. An AAA occurs when the large blood vessel that supplies blood to the abdomen, pelvis and legs becomes abnormally large or balloons outward. A surgical intervention may be needed to avoid rupture, which is associated with a mortality rate of 80 %. While surgery can fix the AAA, it is a serious undertaking; mortality rates during surgery (in the United States) are in the range of 2 % to 6 % for repair under elective, nonemergency, circumstances. Possible complications from surgery include bleeding, infection and kidney or bowel damage. In addition, because coronary artery disease is so common among patients with AAA, a major worry is the risk of post-operative heart trouble. Surgery places a significant strain on the heart and can cause problems such as angina or heart attack.

Hence the difficult decision: should we immediately fix the AAA or should we watchfully wait until the AAA has grown to a size that represents a greater threat to the patient's life? This type of decision is typically made by trained experts, using both statistical data, as well as memories of previous experiences.

A possible set of causal assumptions about the problem is represented in the graph of Figure 8.2. The graph is, once again, confined to a set of three measurement times: ($T_0 < T_1 < T_2$). Let variable X_i (for $i = 0, \ldots, 2$) represent relevant patient information, including the size of the aneurysm at T_i. Let variable A_i indicate, at time T_i, whether the doctors have decided to surgically intervene or to 'watchfully wait'. Let variable W_0 represent unrecorded

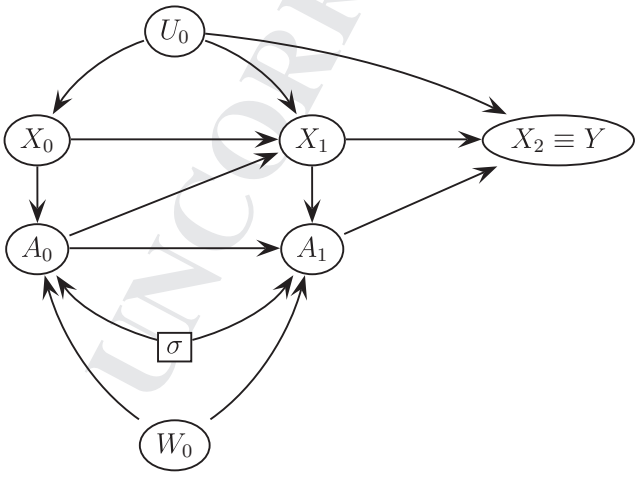

Figure 8.2 Graphical representation of the AAA surveillance example.

factors that vary with the 'doctor in charge' and the treatment centre, whereas U_0 represents unmeasured risk factors, whether of genetic nature or dependent on exposure history.

Two important assumptions are represented in Figure 8.2. The first assumption, corresponding to the missing $U_0 \rightarrow A_k$ arrows, states that no information about the risk factors U_0, unmeasured in the data, has been consulted for the therapeutic interventions $(A_0, A_1, \ldots)$ performed on the study patients. Again, this assumption must be checked by carefully interviewing the doctors in charge. The second assumption, corresponding to the missing $W_0 \rightarrow X_k$ arrows, asserts that the impact of a therapy plan is the same, whoever the doctor, or Health Centre, in charge of the treatment.

It is straightforward to verify, via the usual conditional independence semantics of directed acyclic graphs, that the stability condition (8.5) is satisfied in this particular example. This means that, under the discussed assumptions, observational data on patients under follow-up for AAA may be used as a basis for future therapeutic decisions.

8.10 Statistical inference and computation

Once we have made sure the stability and positivity conditions are valid in the application of interest, the task will be to carry out the actual process of estimation of the causal parameters of interest from the data. One possibility is to impose some form of smoothness across the distributions $p(x_k \mid \overline{x}_{k-1}, \overline{a}_{k1})$ or give them a parametric form (Berzuini and Allemani, 2004). The unknown parameters can be formally included in U_0. The preceding chapter suggests a Bayesian approach to the problem. In a similar vein, we may put a prior distribution on the parameters and then update it into a posterior by using Markov chain Monte Carlo (MCMC) methods to integrate with respect to uncertainty about the estimated parameters, to the stochastic component of the treatment strategy and to the uncertainty in the outcomes.

Then the Markov chain is run on a huge graphical model. To visualize this model, imagine that each sample individual is represented by a corresponding graph, of the type in Figure 8.1 or 8.2, all the resulting graphs being then interconnected through sharing the same model parameters. Upon convergence, the chain will generate samples of the parameters, as if obtained by the correct posterior. Under stability, each $p(x_i \mid \overline{x}_{i-1}, \overline{a}_{i-1})$ distribution is the same under the two regimes, observational and experimental. It follows that the same model, and what we have learned about its parameters, can be used to determine a predictive distribution for Y in a future (hypothetical) patient subject to a therapy strategy e of interest. If we come up with a loss function, $L(Y)$, we can then use the generated MCMC samples to estimate the expectation $E\{L(Y), e\}$, that is the *expected loss* associated with the therapeutic strategy of interest. This will allow us to compare alternative strategies.

We have previously discussed the positivity-related requirements for the method. Are they *always* required? Once we make parametric or smoothness assumptions for the $p(x_i \mid \overline{x}_{i-1}, \overline{a}_{i-1})$ distributions, positivity could be relaxed. In a situation of data scarcity, even under positivity, some of the probabilities might be so small that we are unable to estimate them well on the basis of the available observational data, and this may create problems of convergence of the Markov chain. Smoothing may attenuate this problem.

Another possibility is to adopt a method consisting of two stages (Dawid and Didelez, 2010). In the first stage, given enough data collected under o, we estimate the relevant distributions $p(x_i \mid \overline{x}_{i-1}, \overline{a}_{i-1}\,; o)$ $(i = 0, \ldots, N+1)$. Under stability and positivity, this will also give us the distributions $p(x_i \mid \overline{x}_{i-1}, \overline{a}_{i-1}\,; e)$, for $(i = 0, \ldots, N+1)$. These distributions,

together with the a priori known distributions $p(a_i \mid \overline{x}_i, \overline{a}_{i-1}\,;\, e)$, constitute the main ingredient of the second stage of the procedure, described in the following.

Let h denote a partial history of the form $(\overline{x}_i, \overline{a}_{i-1})$ or $(\overline{x}_i, \overline{a}_i)$ $(0 \leq i \leq N)$. We also include the 'null' history $\emptyset$ and 'full' histories $(\overline{x}_N, \overline{a}_N, y)$. We denote the set of all partial histories by $\mathcal{H}$. Under the experimental regime, e, define a function f on $\mathcal{H}$ by

$$f(h) := E\{L(Y) \mid h\,;\, e\} \tag{8.5}$$

In a medical treatment context, the quantity $f(h)$ represents the expected loss (a negative measure of the overall therapeutic success) for a patient, given his/her past history h. The term 'history' must here be interpreted to denote past evolution of both the disease and the treatment.

Simple application of the laws of probability yields

$$f(\overline{l}_i, \overline{a}_{i-1}) = \sum_{a_i} p(a_i \mid \overline{l}_i, \overline{a}_{i-1}\,;\, e) \times f(\overline{l}_i, \overline{a}_i) \tag{8.6}$$

$$f(\overline{l}_{i-1}, \overline{a}_{i-1}) = \sum_{l_i} p(l_i \mid \overline{l}_{i-1}, \overline{a}_{i-1}\,;\, e) \times f(\overline{l}_i, \overline{a}_{i-1}) \tag{8.7}$$

For h a full history $(\overline{l}_N, \overline{a}_N, y)$, we have $f(h) = L(y)$. Using these as starting values, by successively implementing (8.6) and (8.7) in turn, starting with (8.6) for $i = N + 1$ and ending with (8.7) for $i = 0$, we step down through ever shorter histories until we have computed $f(\emptyset) = E\{L(Y)\,;\emptyset,\ s\}$, the expected loss for strategy e. More generally, we could consider a function Y^* of $(\overline{X}_N, \overline{A}_N, Y)$. Starting now with $f(\overline{x}_N, \overline{a}_N, y) := Y^*(\overline{x}_N, \overline{a}_N, y)$, we can apply the identical steps to arrive at $f(\emptyset) = E\{Y^*\,;\, e\}$, which yields the desired expected loss associated with treatment strategy e. Under suitable further conditions we can combine this recursive method with the selection of an optimal strategy, when it becomes dynamic programming. Although the method is applied when we have data from all possible *static* strategies, it is nevertheless vital that the stability assumption includes all possible dynamic regimes.

When e is a nonrandomized strategy, the distribution of A_i given $\overline{X}_i = \overline{l}_i$, when $\sigma = e$, is degenerate, at $a_i = g_i = g_i(\overline{l}_i\,;\, s)$, say, and the only randomness left is for the variables $(X_0, \ldots, X_N, Y)$. We can now consider $f(h)$ as a function of only the (x_i) appearing in h, since, under e, these then determine the (a_i). Then (8.6) holds automatically, while (8.7) becomes

$$f(\overline{x}_{i-1}) = \sum_{x_i} p(l_i \mid \overline{x}_{i-1}, \overline{g}_{i-1}; e) \times f(\overline{x}_i) \tag{8.8}$$

When, further, the regime e is static, each g_i in the above expressions reduces to the fixed action a_i^* specified by e.

The conditional distributions in (8.6), (8.7) and (8.5) are undefined when the conditioning event has probability 0 under s. One possibility is to define $f(h)$ in (8.6) to be equal to 0 whenever $p(h\,;\ \emptyset,\ e) = 0$.

Robins and Wasserman (1997) warn about the uncritical use of parametric models for the involved conditional distributions, on the grounds that they can lead to a so-called *null*

paradox, which prevents discovering that different plans have the same effect. In the light of these considerations, these authors propose the use of *marginal* or *nested structural models* (Robins, 1998, 2004) that avoid the null paradox.

Computational aspects of the theory and their complex relationship with substantive considerations are illustrated in Chapter 17 in this volume.

8.11 Transparent actions

Let us use the term *transparent actions* to denote the situation where the actions performed under e are only influenced by earlier variables that are observed under both regimes. This is what occurs, for example, in Figure 8.1. Under this assumption, the distribution of Y under the experimental regime can be written as

$$P(Y \mid \sigma = e) = \sum_{\overline{a}_N, \overline{x}_N} \left[\prod_{k=0}^{N} P(a_k \mid pa_{A_k}; \sigma = e) \times \left(\sum_{\overline{u}_N} P(Y \mid pa_Y) \prod_{k=0}^{N} P(x_k \mid pa_{X_k}) \prod_{k=0}^{N} P(u_k) \right) \right] \tag{8.9}$$

where the outer summation is over all possible multivariate configurations of $\overline{a}_N$ and $\overline{x}_N$. The expression on the right-hand side is obtained by factorizing the joint distribution over the graph into a product of conditional distributions of each variable given its parents in the graph, and by then averaging with respect to $(\overline{u}_N, \overline{a}_N, \overline{x}_N)$. A consequence of the transparent action assumption is that the parental sets for the action variables, pa_{A_i}, do not contain u variables, and this justifies the fact that, in the above expression, the $P(a_i \mid pa_{A_i}; \sigma = e)$ factors are not averaged over $\overline{u}_N$.

Importantly, Tian (2008) notes that the expression within square brackets corresponds to the distribution of $(\overline{X}_N, Y)$ under an atomic intervention on $\overline{A}_N$. Let this distribution be shorthanded as $P(Y, \overline{X}_N \mid \sigma = \overline{A}_N)$. Equation (8.9) can then be rewritten as

$$P(Y \mid \sigma = e) = \sum_{\overline{a}_N, \overline{x}_N} \left[P(Y, \overline{x}_N \mid \sigma = \overline{a}_N) \prod_{k=0}^{N} P(a_k \mid pa_{A_k}; \sigma = e) \right] \tag{8.10}$$

Note that, because the $P(a_k \mid pa_{A_k}; \sigma = e)$ terms are *given*, the only unknown quantities in the above expression are the $P(Y, \overline{x}_N \mid \sigma = \overline{a}_N)$, from which the following criterion follows (Shpitser and Pearl, 2006).

Criterion

Under the transparent actions assumption, the distribution of interest, $P(Y \mid \sigma = e)$, is estimable if we are able to identify the distribution of $(\overline{X}_N, Y)$ under a generic atomic intervention on $\overline{A}_N$ from the data.

A complete algorithm for checking identifiability under atomic interventions is the *do*-calculus (Tian and Pearl, 2003; Shpitser and Pearl, 2006), reviewed in Chapter 6 in this volume.

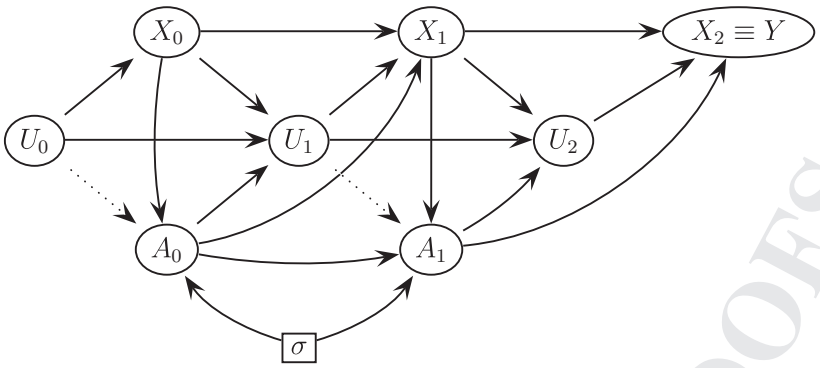

Figure 8.3 Elaboration of the HIV problem of Figure 8.1. Some of the arrows are dotted to indicate that information about the U *variables influenced the action decisions in the observational distribution, but not in the experimental one.*

We conclude that, under the transparent actions assumption, the atomic (static) semantics of the *do*-calculus can be employed to analyse the identifiability of the effects of the dynamic plans.

8.12 Refinements

The above theory can be refined in a number of directions. For example, Tian (2008) considers situations where some X variables are *not* ancestors of Y and refines the theory correspondingly. However, situations in scientific practice where such refinement is crucial have not often been shown.

Another direction is towards problems in which the experimental and the observational distributions are characterized by different graphs, specifically, the former graph being a *subgraph* of the latter. Consider the example of Figure 8.3, which is an elaboration of the HIV graph of Figure 8.1. The elaboration consisted of adding *dotted* arrows from U_k to A_k, (for $k = 0, \ldots, N$). These arrows are intended to be present in the observational graph, *but missing* in the experimental graph. In other words, the graph of Figure 8.3 assumes that, while the data were generated by actions taken in the light of U-information, inferential interest focuses on treatment strategies that do *not* use that information. A possible justification is that the treatments are intended for application in medical routine contexts where expensive genetic monitoring of the virus cannot be afforded. In this case, the distribution of inferential interest, $P(Y \mid \sigma = e)$, is still given by Equation (8.9). This suggests that application of the criterion of the preceding section to the experimental graph (that is to the graph without the dotted arrows) will provide correct guidance in assessing identifiability. In the particular example of Figure 8.3, this will lead to the conclusion that the strategies of interest are identifiable from the data.

8.13 Discussion

Stability and positivity are *sufficient* conditions for the estimability of the effect of treatment strategies in the sense of the present chapter. However, it turns out they are *not* necessary.

The simple stability condition, in particular, requires that, for each time T_i, the conditional distribution of X_i given the earlier observed variables should be the same under both regimes, observational and experimental. This is, indeed, a strong assumption. In many applicative situations, one might not be willing to accept it. Work has been done to establish more general conditions for identifiability (see Dawid and Didelez, 2005, for example), mainly by imposing restrictions on the kinds of information made available for the decisions under the strategy e.

Finally, an important development will be to extend the methods for use when event times, on a continuous support, are explicitly incorporated in the analysis.

Acknowledgements

We thank Dr Silvia Corona for providing insight into the HIV example.

References

Arjas, E. and Parner, J. (2004) Causal reasoning from longitudinal data. *Scandinavian Journal of Statistics*, **31** (2), 171–187.

Berzuini, C. and Allemani, C. (2004) Effectiveness of potent antiretroviral therapy on progression of human immunodeficiency virus: Bayesian modelling and model checking via counterfactual replicates. *Journal of the Royal Statistical Society. Series C (Applied Statistics)*, **53** (4), 633–650.

Dawid, A.P. (1979) Conditional independence in statistical theory (with discussion). **41**, 1–31.

Dawid, A.P. (2002) Influence diagrams for causal modelling and inference. *International Statistical Review*, **70**, 161–189. Corrigenda, **70**, 437.

Dawid, A.P. (2003) Causal inference using influence diagrams: the problem of partial compliance (with discussion), in *Highly Structured Stochastic Systems* (eds P.J. Green, N.L. Hjort and S. Richardson), Oxford University Press, pp. 45–81.

Dawid, A.P. (2010) Beware of the dog!, in *Proceedings of the NIPS 2008 Workshop on Causality. Journal of Machine Learning Research Workshop and Conference Proceedings* (eds D. Janzing, I. Guyon and B. Schoelkopf), vol. 6, pp. 59–86.

Dawid, A.P. and Didelez, V. (2010) Identifying the consequences of dynamic treatment strategies: a decision-theoretic overview. *Statistics Surveys*, **4**, 184–231.

Pearl, J. (2009) *Causality,* 2nd edn. Cambridge: Cambridge University Press.

Pearl, J. and Robins, J. (1995) *Probabilistic Evaluation of Sequential Plans from Causal Models with Hidden Variables*. San Francisco, CA.

Robins, J.M. (1986) A new approach to causal inference in mortality studies with sustained exposure period: application to control of the healthy worker survivor effect. *Mathematical Modelling*, **7**, 1393–1512.

Robins, J.M. (1987) Addendum to a new approach to causal inference in mortality studies with sustained exposure period: application to control of the healthy worker survivor effect. *Computers and Mathematics with Applications*, **14**, 923–945.

Robins, J.M. (1989) The analysis of randomized and nonrandomized aids treatment trials using a new approach to causal inference in longitudinal studies, in *Health Service Research Methodology: A Focus on AIDS* (eds H. Freeman L. Sechrest and A. Mulley), US Public Health Service, pp. 113–159.

Robins, J.M. (1992) Estimation of the time-dependent accelerated failure time model in the presence of confounding factors. *Biometrika*, **79**, 321–324.

Robins, J.M. (1997) Causal inference from complex longitudinal data, in *Latent Variable Modeling and Applications to Causality* (ed. M. Berkane), vol. 120 of Lecture Notes in Statistics, New York: Springer-Verlag, p. 69117.

Robins, J.M. (1998) Structural nested failure time models, in *Survival Analysis*, vol. 6 of *Encyclopedia of Biostatistics* (eds P.K. Andersen and N. Keiding). Chichester, UK: John Wiley & Sons, Ltd. pp. 4372–4389.

Robins, J.M. (2000) Robust estimation in sequentially ignorable missing data and causal inference models, in *Proceedings of the American Statistical Association Section on Bayesian Statistical Science 1999*, vol. 79, p. 610.

Robins, J.M. (2004) Optimal structural nested models for optimal sequential decisions, in *Proceedings of the Second Seattle Symposium on Biostatistics* (eds D.Y. Lin and P. Heagerty). Springer, New York: Springer, p. 189326.

Robins, J.M. and Wasserman, L.A. (1997) Estimation of effects of sequential treatments by reparameterizing directed acyclic graphs, in *Proceedings of the 13th Annual Conference on Uncertainty in Artificial Intelligence* (eds Gerger and P. Shenoy). San Francisco, CA: Morgan Kaufmann Publishers, Francisco, pp. 409–420.

Robins, J.M. Greenland, S. and Hu, F.C. (1999) Estimation of the causal effect of a time-varying exposure on the marginal mean of a repeated binary outcome. *Journal of the American Statistical Association*, **94**, 687–700.

Shpitser, I. and Pearl, J. (2006) Identification of joint interventional distributions in recursive semi-Markovian causal models, in *Proceedings of the Twenty-First National Conference on Artificial Intelligence*, Menlo Park, CA, July 2006.

Tian, J. (2008) Identifying dynamic sequential plans, in *Proceedings of the Uncertainty in Artificial Intelligence Conference,* Helsinki.

Tian, J. and Pearl, J. (2003) *On the Identification of Causal Effects*. Technical Report R-290-L, Department of Computer Science, University of California, Los Angeles.