# Identifying the consequences of dynamic treatment strategies: A decision-theoretic overview* 

A. Philip Dawid<br>Centre for Mathematical Sciences<br>University of Cambridge<br>Wilberforce Road<br>Cambridge CB3 0WB<br>UK<br>e-mail: apd@statslab.cam.ac.uk<br>url: tinyurl.com/2maycn<br>and<br>Vanessa Didelez<br>Department of Mathematics<br>University of Bristol<br>University Walk<br>Bristol BS8 1TW<br>UK<br>e-mail: vanessa.didelez@bristol.ac.uk<br>url: tinyurl.com/2uuteo8


#### Abstract

We consider the problem of learning about and comparing the consequences of dynamic treatment strategies on the basis of observational data. We formulate this within a probabilistic decision-theoretic framework. Our approach is compared with related work by Robins and others: in particular, we show how Robins's ' $G$-computation' algorithm arises naturally from this decision-theoretic perspective. Careful attention is paid to the mathematical and substantive conditions required to justify the use of this formula. These conditions revolve around a property we term stability, which relates the probabilistic behaviours of observational and interventional regimes. We show how an assumption of 'sequential randomization' (or 'no unmeasured confounders'), or an alternative assumption of 'sequential irrelevance', can be used to infer stability. Probabilistic influence diagrams are used to simplify manipulations, and their power and limitations are discussed. We compare our approach with alternative formulations based on causal DAGs or potential response models. We aim to show that formulating the problem of assessing dynamic treatment strategies as a problem of decision analysis brings clarity, simplicity and generality.


AMS 2000 subject classifications: Primary 62C05; secondary 62A01.
Keywords and phrases: Causal inference, $G$-computation, influence diagram, observational study, potential response, sequential decision theory, stability.

Received October 2010.

[^0]
## Contents

1 Introduction ..... 186
1.1 Conditional independence ..... 187
1.2 Overview ..... 188
2 A multistage decision problem ..... 188
3 Regimes and consequences ..... 190
3.1 Inference across regimes ..... 191
4 Evaluation of consequences ..... 191
5 Identifying the ingredients ..... 193
5.1 Control strategies ..... 193
5.2 Stability ..... 193
5.2.1 Some comments ..... 195
5.2.2 Positivity ..... 196
$5.3 \quad G$-recursion ..... 196
6 Extended stability ..... 197
6.1 Preliminaries ..... 198
6.2 Stability regained ..... 199
6.2.1 Sequential randomization ..... 200
6.2.2 Sequential irrelevance ..... 201
7 Influence diagrams ..... 203
7.1 Semantics ..... 203
7.2 Extended stability ..... 204
7.2.1 Sequential randomization ..... 204
7.2.2 Sequential irrelevance ..... 205
7.2.3 Further examples ..... 206
7.2.4 Positivity ..... 207
8 A more general approach ..... 207
8.1 $G$-recursion: General conditions ..... 208
8.2 Extended stability ..... 210
8.2.1 Graphical check ..... 211
8.3 Examples ..... 212
8.3.1 Stability ..... 212
8.3.2 $G$-recursion without stability ..... 213
9 Constructing an admissible sequence ..... 215
9.1 Finding a better sequence ..... 217
9.2 Admissible orderings of $\mathcal{A}$ ..... 217
10 Potential response models ..... 218
10.1 Potential responses and stability ..... 218
10.1.1 Connexions ..... 220
10.2 Potential responses without stability ..... 221
10.2.1 Connexions ..... 222
11 Discussion ..... 223
11.1 What has been achieved? ..... 223
11.2 Syntax and semantics ..... 223
11.3 Statistical inference ..... 224
11.4 Optimal dynamic treatment strategies ..... 225
11.5 Complete identifiability ..... 225
11.6 Other problems ..... 226
Acknowledgment ..... 226
A Two lemmas on DAG-separation ..... 226
References ..... 227

## 1. Introduction

Many important practical problems involve sequential decisions, each chosen in the light of the information available at the time, including in particular the observed outcomes of earlier decisions. As an example, consider long-term anticoagulation treatment, as often given after events such as stroke, pulmonary embolism or deep vein thrombosis. The aim is to ensure that the patient's prothrombin time (INR) is within a target range (which may depend on the diagnosis). Patients on this treatment are monitored regularly, and when their INR is outside the target range the dose of anticoagulant is increased or decreased, so that the dose at any given time is a function of the previous INR observations. Despite the availability of limited guidelines for adjusting the dose, the quality of anticoagulation control achieved is often poor (Rosthøj et al., 2006). Another example is the question of when to initiate antiretroviral therapy for an HIV-1-infected patient. The CD4 cell count at which therapy should be started is a central unresolved issue. Preliminary findings indicate that treatment should be initiated when the CD4 cell count drops below a certain level, i.e. treatment should be a function of the patient's previous CD4 count history (Sterne et al., 2009).

In general, any well-specified way of adjusting the choice of the next decision (treatment or dose to administer) in the light of previous information constitutes a dynamic decision (or treatment) strategy. There will typically be an enormous number of strategies that could be thought of. Researchers would like to be able to evaluate and compare these and, ideally, choose a strategy that is optimal according to a suitable criterion (Murphy, 2003). In many applications, such as the examples given above, it is unlikely that we will have access to large random samples of patients treated under each one of the strategies under consideration. At best, the data available will have been gathered in controlled clinical trials, but often we will have to content ourselves with data from uncontrolled observational studies, with, for example, the treatments being selected by doctors according to informal criteria that we do not know. The key question we address in the present paper is: Under what conditions, and how, could the available data be used to evaluate, compare, and hence choose among, the various decision strategies? When a given strategy can be evaluated from available data it will be termed identifiable.

In principle, our problem can be formulated, represented and solved using the machinery of sequential decision theory, including decision trees and influence
diagrams (Oliver and Smith, 1990; Raiffa, 1968) - and this is indeed the approach that we shall take in this paper. However, this machinery does not readily provide us with an answer to the question of when data obtained, for example, from an observational study will be sufficiently informative to identify a given strategy. Here, we shall be concerned only with issues around potential biases in the data, rather than their completeness. Thus wherever necessary we suppose that the quantity of data available is sufficient to estimate, to any desired precision, the parameters of the process that actually produced those data. However, that process might still differ from that in the new decision problem at hand. We shall therefore propose simple and empirically meaningful conditions (which can thus be meaningfully criticised) under which it is appropriate and possible to make use of the available parameter estimates, and we shall develop formulae for doing this. These conditions will be termed stability due to the way they relate observational and interventional regimes. We shall further discuss how one might justify this stability condition by including unobservable variables into the decision theoretic framework, and by using influence diagrams.

Our proposal is closely related to the seminal work of Robins (Robins, 1986, 1987, 1989, 1997). Much of Robins (1986) takes an essentially decision theoretic approach, while also using the framework of structured tree graphs as well as potential responses (and later using causal direct acyclic graphs (DAGs), see Robins (1997)). He shows that under conditions linking hypothetical studies, where the different treatment strategies to be compared are applied, identifiability can be achieved. Robins calls these conditions sequential randomization (and later no unmeasured confounding, see e.g. Robins (1992)). While these are often formalised using potential responses, a closer inspection of Robins (1986) (or especially Robins (1997)) reveals that all that is needed is an equality of conditional distributions under different regimes, which is what our stability conditions state explicitly. Furthermore, Robins (1986) introduces the $G$-computation algorithm as a method to evaluate a sequential strategy, and contrasts it with traditional regression approaches that yield biased results even when stability or sequential randomization holds (Robins, 1992). We shall demonstrate below that, assuming stability, this $G$-computation algorithm arises naturally out of our decision-theoretic analysis, where it can be recognized as a version of the fundamental 'backward induction' recursion algorithm of dynamic programming.

### 1.1. Conditional independence

The technical underpinning for our decision-theoretic formulation is the application of the language and calculus of conditional independence (Dawid, 1979, 2002) to relate observable variables of two types: 'random' variables and 'decision' (or 'intervention') variables. This formalism is used to express relationships that may be assumed between the probabilistic behaviour of random variables under differing regimes (e.g., observational and interventional). Nevertheless, although it does greatly clarify and simplify analysis, this particular language is not indispensable: everything we do could, if so desired, be expressed directly
in terms of relationships between probability distributions for observable variables. Thus no essential additional ingredients are being added to the standard formulation of statistical decision theory.

In many cases the conditional independence relations we work with can be represented by means of a graphical display: the influence diagram (ID). Once again, although enormously helpful this is, in a formal sense, only an optional extra. Moreover, although we pay special attention to problems that can be represented by influence diagrams, there are yet others, still falling under our general approach, where this is not possible.

Inessential though these ingredients are, we nevertheless suggest that it is well worth the effort of mastering the basic language and properties, both algebraic and graphical, of conditional independence. In particular, these allow very simple derivations of the logical consequences of assumptions made (Dawid, 1979; Lauritzen et al., 1990).

### 1.2. Overview

In $\S \S 2$ and 3 we set out the basic ingredients of our problem and our notation. Section 4 identifies a simple recursion that can be used to calculate the consequence of applying a given treatment regime when the appropriate probabilistic ingredients are available. In $\S 5$ we consider how these ingredients might be come by, and show that the simple stability condition mentioned above allows estimation of these ingredients - and thus, by application of the procedure of $G$ recursion, of the overall consequence. In $\S \S 6$ and 7 we consider how one might justify this stability condition, starting from a position ('extended stability') that might sometimes be more defensible, and relate various sets of sufficient conditions for this to properties of influence diagrams. Section 8 develops more general conditions, similar to Robins (1987) and Robins (1997), under which $G$-recursion can be justified, while $\S 9$ addresses the question of finding an ordering of the involved variables suitable to carry out $G$-recursion. Finally $\S 10$ shows how analyses based on the alternative formalism of potential responses can be related mathematically to our own development.

## 2. A multistage decision problem

We are concerned with a sequential data-gathering and decision-making process, progressing through a discrete sequence of stages. The archetypical context is that of a sequence of medical treatments applied to a patient over time, each taking into account any interim responses or adverse reactions to earlier treatments, such as the anticoagulation treatment for stroke patients or the decision of when to start antiretroviral therapy for HIV patients. We shall sometimes use this language.

Associated with each patient are two sets of variables: $\mathcal{L}$, the set of observable variables, and $\mathcal{A}$, the set of action variables. The variables in $\mathcal{A}$ can, in principle, be manipulated by external intervention, while those in $\mathcal{L}$ are generated and
revealed by Nature. The variables in $\mathcal{L} \cup \mathcal{A}$ are termed domain variables. There is a distinguished variable $Y \in \mathcal{L}$, the response variable, of special concern.

A specified sequence $\mathcal{I}:=\left(L_{1}, A_{1}, \ldots, L_{N}, A_{N}, L_{N+1} \equiv Y\right)$, where $A_{i} \in \mathcal{A}$ and the $L_{i}$ are disjoint subsets of $\mathcal{L}$, defines the information base. The interpretation is that the variables arise or are observed in that order; $L_{i}$ represents (possibly multivariate, generally time-dependent) patient characteristics or other variables over which we have no control, observable between times $i-1$ and $i ; A_{i}$ describes the treatment action applied to the patient at time $i$; and $Y$ is the final 'response variable' of primary interest.

For simplicity we suppose throughout that all these variables exist and can be observed for every patient. Thus we do not directly consider cases where, e.g., $Y$ is time to death, which might occur before some of the $L$ 's and $A$ 's have had a chance to materialize. However our analyses could readily be elaborated to handle such extensions.

When the aim is to control $Y$ through appropriate choices for the action variables $\left(A_{i}\right)$, any principled approach will involve making comparisons, formal or informal, between the implied distributions of $Y$ under a variety of possible strategies for choosing the $\left(A_{i}\right)$. For example, we might have specified a loss $L(y)$ associated with each outcome $y$ of $Y$, and desire to minimise its expectation $\mathrm{E}\{L(Y)\} .{ }^{1}$ Any such decision problem can be solved as soon as we know the relevant distributions for $Y$ (Dawid, 2000, Section 6).

The simplest kind of strategy is to apply some fixed pre-defined sequence of actions, irrespective of any observations on the patient: we call this a static or unconditional strategy (Pearl (2009) terms it atomic). However in realistic contexts static strategies, which do not take any account of accruing information, will be of little interest. In particular, under a decision-theoretically optimal strategy the action to be taken at any stage must typically be chosen to respond appropriately to the data available at that stage (Murphy, 2003; Robins, 1989).

A non-randomized dynamic treatment strategy (with respect to a given information base $\mathcal{I}$ ) is a rule that determines, for each stage $i$ and each configuration (or partial history) $h_{i}:=\left(l_{1}, a_{1}, \ldots, a_{i-1}, l_{i}\right)$ for the variables $\left(L_{1}, A_{1}, \ldots\right.$, $A_{i-1}, L_{i}$ ) available prior to that stage, the value $a_{i}$ of $A_{i}$ that is then to be applied.

Any decision-theoretically optimal strategy can always be chosen to be nonrandomized. Nevertheless, for added generality we shall also consider randomized ${ }^{2}$ dynamic treatment strategies. Such a strategy determines, for each stage $i$ and associated partial history $h_{i}$, a probability distribution for $A_{i}$, describing the random way in which the next action $A_{i}$ is to be generated. When every such randomization distribution is degenerate at a single action this reduces to a non-randomized strategy.

[^1]Suppose now we wish to compare a number of such strategies. If we knew or could estimate the full probabilistic structure of all the variables under each of these, we could simply calculate and compare directly the various distributions for the response $Y$. As outlined in the introduction, our principal concern in this paper is how to obtain such distributional knowledge, when in many cases the only data available will have been gathered under purely observational or other circumstances that might be very different from the strategies we want to compare. To clarify the potential difficulties, consider a statistician or scientist S, who has obtained data on a collection of variables for a large number of patients. She wishes to use her data, if possible, to identify and compare the consequences of various treatment interventions or policies that might be contemplated for some new patient. A major complication, and the motivation for much work in this area, is that S's observational data will often be subject to 'confounding'. For example, S's observations may include actions $\left(A_{i}\right)$ that have been determined by a doctor D , partly on the basis of additional private information D has about the patient, over and above the variables $S$ has measured. Then knowledge of the fact that D has selected an act $A_{i}=a_{i}$, by virtue of that being correlated with unobserved private information D has that may also be predictive of the response $Y$, could affect the distribution of $Y$ in this observational regime in a way different from what would occur if D had no such private information, or if S had herself chosen the value of $A_{i}$. In particular, without giving careful thought to the matter we cannot simply assume that probabilistic behaviour seen under the observational regime will be directly relevant to other, e.g. interventional, regimes of interest.

## 3. Regimes and consequences

In general, we consider the distribution of all the variables in the problem under a variety of different regimes, possibly but not necessarily involving external intervention. For example, these might describe different locations, time-periods, or contexts in which observations can be made. For simplicity we suppose that the domain variables are the same for all regimes. Formally, we introduce a regime indicator, $\sigma$, taking values in some set $\mathcal{S}$, which specifies which regime is under consideration - and thus which (known or unknown) joint distribution over the domain variables $\mathcal{L} \cup \mathcal{A}$ is operating. Thus $\sigma$ has the logical status of a parameter or decision variable, rather than a random variable. We think of the value $s$ of $\sigma$ as being determined externally, before any observations are made; all probability statements about the domain variables must then be explicitly or implicitly conditional on the value of $\sigma$. We use e.g. $p(y \mid x ; s)$ to denote the conditional density for $Y$, at $y$, given $X=x$, under regime $\sigma=s$. In order to side-step measure-theoretic subtleties, we shall confine attention to the case that all variables considered are discrete; in particular, the terms 'distribution' or 'density' should be interpreted as denoting a probability mass function. However, the basic logic of our arguments does extend to more general cases (albeit with some non-trivial technical complications to handle null events.)

If we know $p(y ; s)$ for all $y$, we can determine, for any function $k(\cdot)$, the expectation $\mathrm{E}\{k(Y) ; s\}$. Often we shall be interested in one or a small number of such functions, e.g. a loss function $k(y) \equiv L(y)$. For definiteness we henceforth consider a fixed given function $k(Y)$, and use the term consequence of $s$ to denote the expectation $\mathrm{E}\{k(Y) ; s\}$ of $k(Y)$ when regime $s$ is followed.

More generally we might wish to focus attention on a subgroup (typically defined in terms of the pre-treatment information $L_{1}$ ), and compare the various 'conditional consequences', given membership of the subgroup. Although we do not address this directly here, it is straightforward to extend our unconditional analysis to this case.

### 3.1. Inference across regimes

In the most usual and useful situation, $\mathcal{S}=\{o\} \cup \mathcal{S}^{*}$, where $o$ is a particular observational regime under which data have been gathered, and $\mathcal{S}^{*}$ is a collection of contemplated interventional strategies with respect to the information base $\left(L_{1}, A_{1}, \ldots, L_{N}, A_{N}, Y\right)$. We wish to use data collected under the observational regime $o$ to identify the consequence of following any of the strategies $e \in \mathcal{S}^{*}$. This means we need to make inference strictly beyond the available data to what would happen, in future cases, under regimes that we have not been able to observe in the past.

It should be obvious, but nonetheless deserves emphasis, that we can not begin to address this problem without assuming some relationships between the probabilistic behaviour of the variables across the differing regimes, both observed and unobserved. Inferences across regimes will typically be highly sensitive to the assumptions made, and the validity of our conclusions will depend on their reasonableness. Although in principle any such assumptions are open to empirical test, using data gathered under all the regimes involved, this will often be impossible in practice. In this case, while it is easy to make assumptions, it can be much harder to justify them. Any justification must involve contextdependent considerations, which we can not begin to address here. Instead we simply aim to understand the logical consequences of making certain assumptions. One message that could be drawn is: if you don't like the consequences, rethink your assumptions.

## 4. Evaluation of consequences

Writing e.g. $\left(L_{1}, L_{2}\right)$ for $L_{1} \cup L_{2}$, we denote $\left(L_{1}, \ldots, L_{i}\right)$ by $\bar{L}_{i}$, with similar conventions for other variables in the problem.

For any fixed regime $s$, we can specify the joint distribution of ( $L_{1}, A_{1}, \ldots, L_{N}$, $\left.A_{N}, Y\right)$, when $\sigma=s$, in terms of its sequential conditional distributions for each variable, given all earlier variables. These comprise:
(i). $p\left(l_{i} \mid \bar{l}_{\underline{i}-1}, \bar{a}_{i-1} ; s\right)$ for $i=1, \ldots, N$.
(ii). $p\left(a_{i} \mid \bar{l}_{i}, \bar{a}_{i-1} ; s\right)$ for $i=1, \ldots, N$.
(iii). $p\left(y \mid \bar{l}_{N}, \bar{a}_{N} ; s\right)$.

Note that (iii) can also be considered as the special case of (i) for $i=N+1$.
With $l_{N+1} \equiv y$, we can factorize the overall joint density as:

$$
\begin{equation*}
p(y, \bar{l}, \bar{a} ; s)=\left\{\prod_{i=1}^{N+1} p\left(l_{i} \mid \bar{l}_{i-1}, \bar{a}_{i-1} ; s\right)\right\} \times\left\{\prod_{i=1}^{N} p\left(a_{i} \mid \bar{l}_{i}, \bar{a}_{i-1} ; s\right)\right\} \tag{1}
\end{equation*}
$$

If we know all the terms in (1), we can simply sum out over all variables but $l_{N+1} \equiv y$ to obtain the desired distribution $p(y ; s)$ of $Y$ under regime $s$, from which we can in turn compute the consequence $\mathrm{E}\{k(Y) ; s\}$.

Alternatively, and more efficiently, this calculation can be implemented recursively, as follows. Let $h$ denote a partial history, of the form $\left(\bar{l}_{i}, \bar{a}_{i-1}\right)$ or $\left(\bar{l}_{i}, \bar{a}_{i}\right)$ $(0 \leq i \leq N)$. We also include the 'null' history $\emptyset$, and 'full' histories ( $\left.\bar{l}_{N}, \bar{a}_{N}, y\right)$. We denote the set of all partial histories by $\mathcal{H}$. Fixing the regime $s$, define a function $f$ on $\mathcal{H}$ by:

$$
\begin{equation*}
f(h):=\mathrm{E}\{k(Y) \mid h ; s\} . \tag{2}
\end{equation*}
$$

Simple application of the laws of probability yields:

$$
\begin{align*}
f\left(\bar{l}_{i}, \bar{a}_{i-1}\right) & =\sum_{a_{i}} p\left(a_{i} \mid \bar{l}_{i}, \bar{a}_{i-1} ; s\right) \times f\left(\bar{l}_{i}, \bar{a}_{i}\right)  \tag{3}\\
f\left(\bar{l}_{i-1}, \bar{a}_{i-1}\right) & =\sum_{l_{i}} p\left(l_{i} \mid \bar{l}_{i-1}, \bar{a}_{i-1} ; s\right) \times f\left(\bar{l}_{i}, \bar{a}_{i-1}\right) . \tag{4}
\end{align*}
$$

For $h$ a full history $\left(\bar{l}_{N}, \bar{a}_{N}, y\right)$, we have $f(h)=k(y)$. Using these as starting values, by successively implementing (3) and (4) in turn, starting with (4) for $i=$ $N+1$ and ending with (4) for $i=1$, we step down through ever shorter histories until we have computed $f(\emptyset)=\mathrm{E}\{k(Y) ; s\}$, the consequence of regime $s .{ }^{3}$

The recursion expressed by (3) and (4) is exactly that underlying the 'extensive form' analysis of sequential decision theory (see e.g. Raiffa (1968)). In particular, under suitable further conditions we can combine this recursive method for evaluation of consequences with the selection of an optimal strategy, when it becomes dynamic programming. This 'step-down histories' approach also applies just as readily to more general probability or decision trees, where the length of the history, and even the variables entering into it, can vary with the path followed. We do not consider such extensions here, but they raise no new issues of principle.

When $s$ is a non-randomized strategy, the distribution of $A_{i}$ given $\bar{L}_{i}=\bar{l}_{i}$, when $\sigma=s$, is degenerate, at $a_{i}=g_{i}=g_{i}\left(\bar{l}_{i} ; s\right)$, say, and the only randomness left is for the variables $\left(L_{1}, \ldots, L_{N}, Y\right)$. We can now consider $f(h)$ as a function of only the $\left(l_{i}\right)$ appearing in $h$, since, under $s$, these then determine the $\left(a_{i}\right)$. Then (3) holds automatically, while (4) becomes:

$$
\begin{equation*}
f\left(\bar{l}_{i-1}\right)=\sum_{l_{i}} p\left(l_{i} \mid \bar{l}_{i-1}, \bar{g}_{i-1} ; s\right) \times f\left(\bar{l}_{i}\right) . \tag{5}
\end{equation*}
$$

[^2]When, further, the regime $s$ is static, each $g_{i}$ in the above expressions reduces to the fixed action $a_{i}^{*}$ specified by $s$.

We remark that the conditional distributions in (i)-(iii) and (2) are undefined when the conditioning event has probability 0 under $s$. The overall results of recursive application of (3) and (4) will not depend on how such ambiguities are resolved. However, for later convenience we henceforth assume that $f(h)$ in (2) is defined as 0 whenever $p(h ; s)=0$. Note that this property is preserved under (3) and (4).

## 5. Identifying the ingredients

In order for the statistician $S$ to be able to apply the above recursive method to calculate the consequence of some contemplated regime $s$, she needs to know all the ingredients (i), (ii) and (iii). How might such knowledge be attained?

### 5.1. Control strategies

Consider first the term $p\left(a_{i} \mid \bar{l}_{i}, \bar{a}_{i-1} ; s\right)$ in (ii), as needed for (3). It will often be the case that for the regimes $s$ of interest this is known a priori to the statistician S for all $i$. For instance we might be interested in strategies for initiating antiretroviral treatment of HIV patients as soon as the CD4 count has dropped below a given value $c$. The strategy therefore fully determines the value of the binary $A_{i}$ given the previous covariate history $\bar{l}_{i}$ as long as this includes information on the CD4 counts. In such a case we shall call $s$ a control strategy (with respect to the information base $\mathcal{I}=\left(L_{1}, A_{1}, \ldots, L_{N}, A_{N}, Y\right)$ ). In particular this will typically be the case when $s$ is a (possibly randomized) dynamic strategy, as introduced in $\S 2$.

### 5.2. Stability

More problematic is the source of knowledge of the conditional density $p\left(l_{i} \mid\right.$ $\bar{l}_{i-1}, \bar{a}_{i-1} ; s$ ) in (i) as required for (4) (including, as a special case, that of $p\left(y \mid \bar{l}_{N}, \bar{a}_{N} ; s\right)$ in (iii)).

If we observed many instances of regime $s$, we may be able to estimate this directly; but typically we will be interested in assessing the consequences of various contemplated regimes (e.g. control strategies) that we have never yet observed. The problem then becomes: under what conditions can we use probability distributions assessed under one regime to deduce the required conditional probabilities, (i) and (iii), under another?

In the application of most interest, we have $\mathcal{S}=\{o\} \cup \mathcal{S}^{*}$, where $o$ is an observational regime under which data have been gathered, and $\mathcal{S}^{*}$ is a collection of contemplated interventional strategies. If we can use data collected under the observational regime $o$ to identify the consequence of following any of the strategies $e \in \mathcal{S}^{*}$, we will be in a position to compare the consequences of
different interventional strategies (and thus, if desired, choose an optimal one) on the basis of data collected in the single regime $o$.

In general, the distribution of $L_{i}$ given $\left(\bar{L}_{i-1}, \bar{A}_{i-1}\right)$ will depend on which regime is in operation. Even application of a control strategy might well have effects on the joint distribution of all the variables, beyond the behaviour it directly specifies for the actions. For example, consider an educational experiment in which we can select certain pupils to undergo additional home tutoring. Such an intervention can not be imposed without subjecting the pupil and his family to additional procedures and expectations, which would probably be different if the decision to undergo extra tutoring had come directly from the pupil, and possibly different again if it had come from the parents. Consequently we can not necessarily assume that the distribution of $L_{i}$ given $\left(\bar{L}_{i-1}, \bar{A}_{i-1}\right)$ assessed under the observational regime will be the same as that for an interventional strategy, or that it would be the same for different interventional strategies.

It will clearly be helpful when we can impose this assumption - and so be able to identify the required interventional distributions of $L_{i}$ given $\left(\bar{L}_{i-1}, \bar{A}_{i-1}\right)$ with those assessed under the observational regime. We formalize this assumption as follows:

Definition 5.1 We say that the problem exhibits simple stability, with respect to the information base $\mathcal{I}=\left(L_{1}, A_{1}, \ldots, L_{N}, A_{N}, Y\right)$ and the set $\mathcal{S}$ of regimes if, with $\sigma$ denoting the non-random regime indicator taking values in $\mathcal{S}$ :

$$
\begin{equation*}
L_{i} \Perp \sigma \mid\left(\bar{L}_{i-1}, \bar{A}_{i-1}\right) \quad(i=1, \ldots, N+1) . \tag{6}
\end{equation*}
$$

Here and throughout, we use the notation and theory of conditional independence introduced by Dawid (1979), as generalized as in Dawid (2002) to apply also to problems involving decision or parameter variables. In words, condition (6) asserts that the stochastic way in which $L_{i}$ arises, given the previous values of the $L$ 's and $A$ 's, should be the same, irrespective of which regime in $\mathcal{S}$ is in operation. More precisely, expressed in terms of densities, (6) requires that, for each $i=1, \ldots, N+1$, there exist some common conditional density specification $q\left(L_{i}=l_{i} \mid \bar{L}_{i-1}=\bar{l}_{i-1}, \bar{A}_{i-1}=\bar{a}_{i-1}\right)$ such that, for each $s \in \mathcal{S}$,
$p\left(L_{i}=l_{i} \mid \bar{L}_{i-1}=\bar{l}_{i-1}, \bar{A}_{i-1}=\bar{a}_{i-1} ; s\right)=q\left(L_{i}=l_{i} \mid \bar{L}_{i-1}=\bar{l}_{i-1}, \bar{A}_{i-1}=\bar{a}_{i-1}\right)$
whenever the conditioning event has positive probability under regime $s$.
As will be described further in $\S 7$ below, it is often helpful (though never essential) to represent conditional independence properties graphically, using the formalism of influence diagrams (IDs): such diagrams have very specific semantics, and can facilitate logical arguments by displaying implied properties in a particularly transparent form (Dawid, 2002). The appropriate graphical encoding of property (6) for $i=1,2$ and 3 is shown in Figure 1. The specific property (6) is represented by the absence of arrows from $\sigma$ to $L_{1}, L_{2}$, and $Y \equiv L_{3}$. For general $N$ we simply supplement the complete directed graph on $\left(L_{1}, A_{1}, \ldots, L_{N}, A_{N}, Y\right)$ with an additional regime node $\sigma$, and an arrow from $\sigma$ to each $A_{i}$.


FIG 1. Influence diagram: stability.

### 5.2.1. Some comments

An important question is how we should assess whether property (6) holds in any given situation. It could in principle be tested empirically, if we could collect data under all regimes. In practice this is usually impossible, and other arguments for or against its appropriateness would be brought to bear. Whether or not the simple stability property can be regarded as appropriate in any application will depend on the overall context of the problem. In particular, it will depend on the specific information base involved. For example, if $e$ is a control strategy with respect to S's information base, and $o$ an observational regime under which the doctor D chooses the $\left(A_{i}\right)$ on the basis of private information not represented in S's information base, possibly associated with $L_{i}$, then, for $\mathcal{S}=\{o, e\}$, we might well expect (6) to be violated. This is often described as (potential) confounding.

The simple stability property (6) is our version of a condition termed 'sequential randomization' (Robins, 1986, 1997) or 'no unmeasured confounding' (Robins, 1992; Robins, Hernán and Brumback, 2000) or 'sequential ignorability' (Robins, 2000). The connexions become particularly clear when comparing (6) with the equalities derived in Theorem 3.1 of Robins (1997), which we consider in more detail in $\S 10.1 .1$ below. These alternative names suggest particular situations where stability should be satisfied, such as when the data have been gathered under an observational regime where the actions were indeed physically sequentially randomized; or when S's information base contains all the information the doctor D has used in choosing the $\left(A_{i}\right)$. However, we emphasise that our property (6) can be meaningfully considered even without referring to any 'potential confounder' variables; and that if (as in $\S 6$ below) we do choose to introduce such further variables to help us assess whether (6) holds, nevertheless the property itself must hold or fail quite independently of which additional variables (if any) are considered.

In any case, because stability is a property of the relationship between different regimes, it can never be empirically established on the basis of data collected under only one (e.g., observational) regime, nor can it be deduced from properties assumed to hold for just one such regime.

### 5.2.2. Positivity

The purpose of invoking simple stability (with respect to $\mathcal{S}=\{o\} \cup \mathcal{S}^{*}$ ) is to get a handle on (4) for an unobserved interventional strategy $s=e \in \mathcal{S}^{*}$, using data obtained in the observational regime $o$. Intuitively, under simple stability we can replace $p\left(l_{i} \mid \bar{l}_{i-1}, \bar{a}_{i-1} ; e\right)$ by $p\left(l_{i} \mid \bar{l}_{i-1}, \bar{a}_{i-1} ; o\right)$, which is estimable from the observational data. However, some care is needed on account of the positivity qualification following (7). If, for example, we want to assess the consequence of a static interventional strategy $e$, which always applies some pre-specified action sequence $\bar{a}^{*}$, we clearly will be unable to do so using data from an observational regime in which the probability of obtaining that particular sequence of actions is zero. (Pragmatically it may still be difficult to do so if that probability is non-zero but so small that we are unable to estimate it well from available observational data. However we ignore that difficulty here, supposing that the data are sufficiently extensive that we can indeed get good estimates of all probabilities under $o$ ).

In order to avoid this problem, we impose the positivity (absolute continuity) condition:

Definition 5.2 We say the problem exhibits positivity if, for any $e \in \mathcal{S}^{*}$, the joint distribution of $\left(\bar{L}_{N}, \bar{A}_{N}, Y\right)$ under $P_{e}$ is absolutely continuous with respect to that under $P_{o}$, i.e.

$$
\begin{equation*}
p(E ; e)>0 \Rightarrow p(E ; o)>0 \tag{8}
\end{equation*}
$$

for any event $E$ defined in terms of $\left(\bar{L}_{N}, \bar{A}_{N}, Y\right)$. We write this as $P_{e} \ll P_{o}$.
In our discrete set-up, it is clearly enough to demand (8) whenever $E$ comprises a single sequence $\left(\bar{l}_{N}, \bar{a}_{N}, y\right)$. Denoting by $\mathcal{O}, \mathcal{E}$ the sets of partial histories having positive probability under, respectively, regimes $o$ and $e$, we can restate (8) as

$$
\begin{equation*}
\mathcal{E} \subseteq \mathcal{O} \tag{9}
\end{equation*}
$$

### 5.3. G-recursion

Let $e \in \mathcal{S}^{*}$. Given enough data collected under $o$ we can identify $p\left(l_{i} \mid \bar{l}_{i-1}, \bar{a}_{i-1} ; o\right)$ $(i=1, \ldots, N+1)$ for $\left(\bar{l}_{i-1}, \bar{a}_{i-1}\right) \in \mathcal{O}$. Under simple stability (7) and positivity (9), this will also give us $p\left(l_{i} \mid \bar{l}_{i-1}, \bar{a}_{i-1} ; e\right)(i=1, \ldots, N+1)$ for all $\left(\bar{l}_{i-1}, \bar{a}_{i-1}\right) \in \mathcal{E}$. If, further, $e$ is a control strategy, then using the known form for $p\left(a_{i} \mid \bar{l}_{i}, \bar{a}_{i-1} ; e\right)\left(\left(\bar{l}_{i}, \bar{a}_{i}\right) \in \mathcal{E}\right)$, we have all the ingredients to apply (3) and (4) and thus identify the consequence of regime $e$ from data collected under $o$.

Specifically, we have

$$
\begin{align*}
f\left(\bar{l}_{i}, \bar{a}_{i-1}\right) & =\sum_{a_{i}} p\left(a_{i} \mid \bar{l}_{i}, \bar{a}_{i-1} ; e\right) \times f\left(\bar{l}_{i}, \bar{a}_{i}\right)  \tag{10}\\
f\left(\bar{l}_{i-1}, \bar{a}_{i-1}\right) & =\sum_{l_{i}} p\left(l_{i} \mid \bar{l}_{i-1}, \bar{a}_{i-1} ; o\right) \times f\left(\bar{l}_{i}, \bar{a}_{i-1}\right) . \tag{11}
\end{align*}
$$

We start the recursion with

$$
f\left(\bar{l}_{N}, \bar{a}_{N}\right) \equiv \mathrm{E}\left\{k(Y) \mid \bar{l}_{N}, \bar{a}_{N} ; e\right\}= \begin{cases}\mathrm{E}\left\{k(Y) \mid \bar{l}_{N}, \bar{a}_{N} ; o\right\} & \text { if }\left(\bar{l}_{N}, \bar{a}_{N}\right) \in \mathcal{E} \\ 0 & \text { otherwise }\end{cases}
$$

(using simple stability for $i=N+1$ ), and exit with the desired interventional consequence $f(\emptyset) \equiv \mathrm{E}\{k(Y) ; e\}$.

We refer to the above method as $G$-recursion. ${ }^{4}$
For the case that $e$ is a non-randomized strategy, $G$-recursion can be based on (5), becoming

$$
\begin{equation*}
f\left(\bar{l}_{i-1}\right)=\sum_{l_{i}} p\left(l_{i} \mid \bar{l}_{i-1}, \bar{g}_{i-1} ; o\right) \times f\left(\bar{l}_{i}\right) \tag{12}
\end{equation*}
$$

starting with $f\left(\bar{l}_{N}\right)=\mathrm{E}\left\{k(Y) \mid \bar{l}_{N}, \bar{g}_{N} ; o\right\}$. The $G$-computation formula (Robins, 1986) is the algebraic formula for $f(\emptyset)$ in terms of $f\left(\bar{l}_{N}\right)$ that results when we write out explicitly the successive substitutions required to perform this recursion.

Finally we remark that, when the simple stability property (6) holds for $\left(L_{1}, A_{1}, \ldots, L_{N}, A_{N}, Y\right)$, it also holds for $\left(L_{1}, A_{1}, \ldots, L_{N}, A_{N}, Y^{*}\right)$, where $Y^{*}$ is any function of $\left(L_{1}, A_{1}, \ldots, L_{N}, A_{N}, Y\right)$. For $i \leq N$ there is nothing new to show, while (6) for $i=N+1$ follows easily for $Y^{*}$ when it holds for $Y$, using general properties of conditional independence (Dawid, 1979). It is also easy to see that when positivity, Definition 5.2 , holds for $\left(\bar{L}_{N}, \bar{A}_{N}, Y\right)$ it likewise holds for $\left(\bar{L}_{N}, \bar{A}_{N}, Y^{*}\right)$. Consequently, under the same conditions that allow $G$-recursion to compute the interventional distribution of $Y$, we can use it to compute that of $Y^{*}$. In particular (see footnote 1), this will allow us to evaluate the expected loss of applying $e$, even when the loss function depends on all of $\left(\bar{L}_{N}, \bar{A}_{N}, Y\right)$.

## 6. Extended stability

We have already alluded to the possibility that, in many applications, the simple stability assumption (6) might not be easy to justify directly. This might be the case, in particular, when we are concerned about the possibility of 'confounding effects' due to unobserved influential variables.

In such a case we might proceed by constructing a more detailed model, incorporating a collection $\mathcal{U}$ of additional, possibly unobserved, variables; and investigate its implications. These unobserved variables might be termed 'sequential (potential) confounders'. Under certain additional assumptions to be discussed below, we might then be able to deduce that simple stability does, after all, apply. This programme can be helpful when the assumptions involving the additional variables are easier to justify than assumptions referring only to

[^3]the variables of direct interest. We here initially express these additional assumptions purely algebraically, in terms of conditional independence; in $\S 7$ we shall conduct a parallel analysis utilising influence diagrams to facilitate the expression and manipulation of the relevant conditional independencies.

Reasoning superficially similar to ours has been conducted by Pearl and Robins (1995) and Robins (1997). However, that is mostly based on the assumed existence of a 'causal DAG' representation of the problem. We once again emphasise that the simple stability property (6) is always meaningful of itself, and its truth or falsity can not rely on the possibility of carrying out such a programme of reduction from a more complex model including unobservable variables.

### 6.1. Preliminaries

We shall specifically investigate models having a property we term extended stability. Such a model again involves a collection $\mathcal{L}$ of observable domain variables (including a response variable $Y$ ) and a collection $\mathcal{A}$ of action domain variables, together with a regime indicator variable $\sigma$ taking values in $\mathcal{S}=\{o\} \cup \mathcal{S}^{*}$. But now we also have the collection $\mathcal{U}$ of unobservable domain variables (for simplicity we suppose throughout that which variables are observed or unobserved is the same under all regimes considered). Let $\mathcal{I}^{\prime}$ denote an ordering of all these observable and unobservable domain variables (typically, though not necessarily, their time-ordering). As before we assume that $A_{i-1}$ comes before $A_{i}$ in this ordering. We term $\mathcal{I}^{\prime}$ an extended information base. Let $L_{i} \subseteq \mathcal{L}$ [resp., $\left.U_{i} \subseteq \mathcal{U}\right]$ denote the set of observed [resp., unobserved] variables between $A_{i-1}$ and $A_{i}$.

Definition 6.1 We say that the problem exhibits extended stability with respect to the extended information base $\mathcal{I}^{\prime}$ and the set $\mathcal{S}$ of regimes if, for $i=1, \ldots, N+1$,

$$
\begin{equation*}
\left(U_{i}, L_{i}\right) \Perp \sigma \mid\left(\bar{U}_{i-1}, \bar{L}_{i-1}, \bar{A}_{i-1}\right) \tag{13}
\end{equation*}
$$

(If the $\left(U_{i}\right)$ were observable, this would be identical with the definition of simple stability.)

Under extended stability the marginal distribution of $U_{1}$ is supposed the same in both regimes, as is the conditional distribution of $U_{2}$ given $\left(U_{1}, L_{1}, A_{1}\right)$, etc. Similarly, the distributions of $L_{1}$ given $U_{1}$, of $L_{2}$ given $\left(U_{1}, L_{1}, A_{1}, U_{2}\right), \ldots$, and finally of $Y\left(=L_{N+1}\right)$ given $\left(U_{1}, L_{1}, A_{1}, \ldots, U_{N}, L_{N}, A_{N}\right)$, are all supposed to be independent of the regime operating.

There is a corresponding extension of Definition 5.2:
Definition 6.2 We say the problem exhibits extended positivity if, for any $e \in$ $\mathcal{S}^{*}, P_{e} \ll P_{o}$ as distributions over $\left(\bar{L}_{N}, \bar{U}_{N}, \bar{A}_{N}, Y\right)$; that is, $p(E ; e)>0 \Rightarrow$ $p(E ; o)>0$ and any event $E$ defined in terms of $\left(\bar{L}_{N}, \bar{U}_{N}, \bar{A}_{N}, Y\right)$.

In many problems, though by no means universally, an extended stability assumption might be regarded as more reasonable and defensible than simple stability - so long as appropriate unobserved variables $\mathcal{U}$ are taken into account. For example, this might be the case if we believed that, in the observational
regime, the actions were chosen by a decision-maker who had been able to observe, in sequence, some or all of the variables in the problem, including possibly the $U$ 's; and was then operating a control strategy with respect to this extended information base, so that, when choosing each action, he was taking account of all previous variables in this extended sequence, but nothing else. But even then, as discussed in $\S 5.2$, the extended stability property is a strong additional assumption, that needs to be justified in any particular problem. And again, because it involves the relationships between distributions under different regimes, it can not be justified on the basis of considerations or findings that apply only to one regime.

Unobservable variables can assist in modelling the observational regime and its relationship with the interventional control regimes under consideration. But, because they are unobserved, they can not form part of the information taken into account by such control regimes. Thus we shall still be concerned with evaluating - using $G$-recursion when possible - a regime $e$ that is a control strategy with respect to the observable information base $\mathcal{I}=\left(L_{1}, A_{1}, \ldots, L_{N}, A_{N}, Y\right)$ as introduced in $\S 5.1$. More specifically, in this more general context we define:

Condition 6.1 (Control strategy) The regime e is a control strategy if, for $i=1, \ldots, N$,

$$
\begin{equation*}
A_{i} \Perp \bar{U}_{i} \mid\left(\bar{L}_{i}, \bar{A}_{i-1} ; e\right) \tag{14}
\end{equation*}
$$

and, in addition, the conditional distribution of $A_{i}$, given $\left(\bar{L}_{i}, \bar{A}_{i-1}\right)$, under regime e, is known to the analyst.
Condition 6.1 expresses the property that, under regime $e$, the randomization distribution or other sources of uncertainty about $A_{i}$, given all earlier variables, does not in fact depend on the earlier unobserved variables; and that this conditional distribution is known. The condition will hold, in particular, in the important common case that, under $e, A_{i}$ is fully specified as a function of previous observables.

### 6.2. Stability regained

When there are unobservables in the problem, the extended positivity property of Definition 6.2 will clearly imply the simple positivity property of Definition 5.2. However, even when extended stability holds, the simple stability property, with respect to the observable information base $\left(L_{1}, A_{1}, \ldots, L_{N}, A_{N}, Y\right)$ from which (as is a pragmatic necessity) we have had to exclude the unobserved variables, will typically fail. But we can sometimes incorporate additional background knowledge, most usefully expressed in terms of conditional independence, to show that it does, after all, hold.

We now describe two sets of additional sufficient (though not necessary) conditions, either of which will, when appropriate, allow us to deduce the simple stability property (6) - and with it, the possibility of applying $G$-recursion (ignoring the unobservable variables), as set out in $\S 5.3$. The results in this sec-
tion can be regarded as extending the analysis of Dawid (2002) § 8.3 (see also Guo and Dawid (2010)) to the sequential setting.

### 6.2.1. Sequential randomization

It has frequently been proposed (e.g., Robins $(1986,1997)$ ) that when, under an observational regime, the actions $\left(A_{i}\right)$ have been physically (sequentially) randomized, then simple stability (6) will hold. Indeed, our concept of simple stability has also been termed 'sequential randomization' (Robins, 1986). However we shall be more specific and restrict the term sequential randomization to the special case that we have extended stability and, in addition, Condition 6.2 below holds. We shall show that these properties are indeed sufficient to imply simple stability - but they are by no means necessary.

So consider now the following condition:

## Condition 6.2

$$
\begin{equation*}
A_{i} \Perp \bar{U}_{i} \mid\left(\bar{L}_{i}, \bar{A}_{i-1} ; \sigma\right) \quad(i=1, \ldots, N) . \tag{15}
\end{equation*}
$$

This is essentially a discrete-time version of Definition 2 (ii) of Arjas and Parner (2004), but with the additional vital requirement that the unobservable variables $\mathcal{U}$ involved already be such as to allow us to assume the extended stability property (13). (Without such an underlying assumption there can be no way of relating different regimes together.)

Condition 6.2 requires that, for each regime, any earlier unobserved variables in the extended information base $\mathcal{I}^{\prime}$ can have no further effect on the distribution of $A_{i}$, once the earlier observed variables are taken into account. This will certainly be the case when, under each regime, treatment assignment, at any stage, is determined by some deterministic or randomizing device that only has the values of those earlier observed variables as inputs. While this will necessarily hold for a control strategy with respect to the observed information base, whether or not it is a reasonable requirement for the observational regime will depend on deeper consideration of the specific context and circumstances. It will typically do so if all information available to and utilised by the decisionmaker (the doctor, for instance) in the observational regime is included in $\bar{L}_{i}$, or, indeed, if the actions $\left(A_{i}\right)$ have been physically randomized within levels of $\left(\bar{L}_{i}, \bar{A}_{i-1}\right)$.

Theorem 6.1 Suppose our model exhibits extended stability. If in addition Condition 6.2 holds, then we shall also have the simple stability property (6).

Proof. Our proof will be based on universal general properties of conditional independence, as described by Dawid $(1979,1998)$.

Let $E_{i}, R_{i}, H_{i}$ denote, respectively, the following assertions:

$$
\left.\left.\begin{array}{rl}
E_{i} & : \\
R_{i} & : \\
\left.L_{i}, U_{i}\right) \Perp \sigma \mid\left(\bar{L}_{i-1}, \bar{U}_{i-1}, \bar{A}_{i-1}\right) \\
H_{i} & : \\
\bar{U}_{i}\left|\left(L_{i}, \bar{U}_{i}\right) \Perp \sigma\right|\left(\bar{A}_{i-1} ; \sigma\right) \\
i-1
\end{array}\right) \bar{A}_{i-1}\right)
$$

Extended stability is equivalent to $E_{i}$ holding for all $i$, so we assume that; while $R_{i}$ is just Condition 6.2 , which we are likewise assuming for all $i$. We shall show that these assumptions imply $H_{i}$ for all $i$, which in turn implies $L_{i} \Perp \sigma \mid\left(\bar{L}_{i-1}, \bar{A}_{i-1}\right)$, i.e., simple stability.

We proceed by induction. Since $E_{1}$ and $H_{1}$ are both equivalent to $\left(L_{1}, U_{1}\right) \Perp \sigma$, $H_{1}$ holds.

Suppose now $H_{i}$ holds. Conditioning on $L_{i}$ yields

$$
\begin{equation*}
\bar{U}_{i} \Perp \sigma \mid\left(\bar{L}_{i}, \bar{A}_{i-1}\right), \tag{16}
\end{equation*}
$$

and this together with $R_{i}$ is equivalent to $\bar{U}_{i} \Perp\left(A_{i}, \sigma\right) \mid\left(\bar{L}_{i}, \bar{A}_{i-1}\right)$, which on conditioning on $A_{i}$ then yields

$$
\begin{equation*}
\bar{U}_{i} \Perp \sigma \mid\left(\bar{L}_{i}, \bar{A}_{i}\right) . \tag{17}
\end{equation*}
$$

Also, by $E_{i+1}$ we have

$$
\begin{equation*}
\left(L_{i+1}, U_{i+1}\right) \Perp \sigma \mid\left(\bar{L}_{i}, \bar{U}_{i}, \bar{A}_{i}\right) . \tag{18}
\end{equation*}
$$

Taken together, (17) and (18) are equivalent to $H_{i+1}$, so the induction is established.

### 6.2.2. Sequential irrelevance

Another possible condition is:
Condition 6.3

$$
\begin{equation*}
L_{i} \Perp \bar{U}_{i-1} \mid\left(\bar{L}_{i-1}, \bar{A}_{i-1} ; \sigma\right) \quad(i=1, \ldots, N+1) . \tag{19}
\end{equation*}
$$

In contrast to (15), (19) does permit the unobserved variables to date, $\bar{U}_{i}$, to influence the next action $A_{i}$ (which can however only happen in the observational regime), as well as the current observable $L_{i}$; but they do not affect the subsequent development of the $L$ 's (including, in particular, the response variable $Y$ ).
Theorem 6.2 Suppose:
(i). Extended stability, (13), holds.
(ii). Sequential irrelevance, Condition 6.3, holds for the observational regime $\sigma=o:$

$$
\begin{equation*}
L_{i} \Perp \bar{U}_{i-1} \mid\left(\bar{L}_{i-1}, \bar{A}_{i-1} ; \sigma=o\right) \quad(i=1, \ldots, N+1) . \tag{20}
\end{equation*}
$$

(iii). Extended positivity, as in Definition 6.2, holds.

Then we shall have simple stability:

$$
\begin{equation*}
L_{i} \Perp \sigma \mid\left(\bar{L}_{i-1}, \bar{A}_{i-1}\right) \quad(i=1, \ldots, N+1) . \tag{21}
\end{equation*}
$$

Moreover, sequential irrelevance holds under any regime:

$$
\begin{equation*}
L_{i} \Perp \bar{U}_{i-1} \mid\left(\bar{L}_{i-1}, \bar{A}_{i-1} ; \sigma\right) \quad(i=1, \ldots, N+1) . \tag{22}
\end{equation*}
$$

Proof. Let $k\left(L_{i}\right)$ be a bounded real function of $L_{i}$, and, for each regime $s \in \mathcal{S}$, let $h\left(\bar{U}_{i-1}, \bar{L}_{i-1}, \bar{A}_{i-1} ; s\right)$ be a version of $\mathrm{E}\left\{k\left(L_{i}\right) \mid \bar{U}_{i-1}, \bar{L}_{i-1}, \bar{A}_{i-1} ; s\right\}$.

By (20) there exists $f\left(\bar{L}_{i-1}, \bar{A}_{i-1}\right)$ such that

$$
\begin{equation*}
h\left(\bar{U}_{i-1}, \bar{L}_{i-1}, \bar{A}_{i-1} ; o\right)=f\left(\bar{L}_{i-1}, \bar{A}_{i-1}\right) \quad \text { a.s. }\left[P_{o}\right] \tag{23}
\end{equation*}
$$

whence, from (8), for all $s \in \mathcal{S}$,

$$
\begin{equation*}
h\left(\bar{U}_{i-1}, \bar{L}_{i-1}, \bar{A}_{i-1} ; o\right)=f\left(\bar{L}_{i-1}, \bar{A}_{i-1}\right) \quad \text { a.s. }\left[P_{s}\right] \tag{24}
\end{equation*}
$$

Also, from (13),

$$
\begin{equation*}
L_{i} \Perp \sigma \mid \bar{U}_{i-1}, \bar{L}_{i-1}, \bar{A}_{i-1} \tag{25}
\end{equation*}
$$

and so there exists $g\left(\bar{U}_{i-1}, \bar{L}_{i-1}, \bar{A}_{i-1}\right)$ such that, for all $s \in \mathcal{S}$,

$$
\begin{equation*}
h\left(\bar{U}_{i-1}, \bar{L}_{i-1}, \bar{A}_{i-1} ; s\right)=g\left(\bar{U}_{i-1}, \bar{L}_{i-1}, \bar{A}_{i-1}\right) \quad \text { a.s. }\left[P_{s}\right] . \tag{26}
\end{equation*}
$$

In particular,

$$
\begin{equation*}
h\left(\bar{U}_{i-1}, \bar{L}_{i-1}, \bar{A}_{i-1} ; o\right)=g\left(\bar{U}_{i-1}, \bar{L}_{i-1}, \bar{A}_{i-1}\right) \quad \text { a.s. }\left[P_{o}\right] \tag{27}
\end{equation*}
$$

so that, again using (8),

$$
\begin{equation*}
h\left(\bar{U}_{i-1}, \bar{L}_{i-1}, \bar{A}_{i-1} ; o\right)=g\left(\bar{U}_{i-1}, \bar{L}_{i-1}, \bar{A}_{i-1}\right) \quad \text { a.s. }\left[P_{s}\right] \tag{28}
\end{equation*}
$$

Combining (24), (26) and (28), we obtain

$$
\begin{equation*}
h\left(\bar{U}_{i-1}, \bar{L}_{i-1}, \bar{A}_{i-1} ; s\right)=f\left(\bar{L}_{i-1}, \bar{A}_{i-1}\right) \quad \text { a.s. }\left[P_{s}\right] \tag{29}
\end{equation*}
$$

Since this property holds for all $s \in \mathcal{S}$ and every bounded real function $k\left(L_{i}\right)$, we deduce

$$
\begin{equation*}
L_{i} \Perp\left(\bar{U}_{i-1}, \sigma\right) \mid\left(\bar{L}_{i-1}, \bar{A}_{i-1}\right) \tag{30}
\end{equation*}
$$

from which both (21) and (22) follow.
It is worth noting that we do not need the full force of extended stability for the above proof, but only (25). In particular, we could allow arbitrary dependence of $U_{i}$ on any earlier variables, including $\sigma$. We note further that the above proof makes essential use of the extended positivity property of Definition 6.2: (21) can not be deduced from extended stability and Condition 6.3 making use of the standard conditional independence axioms (Dawid, 1998, 2001; Pearl and Paz, 1987) alone.

Although we can certainly deduce simple stability when we can assume the conditions of either Theorem 6.1 or Theorem 6.2 , it can also arise our of extended stability in other ways. For example, this can be so when Condition 6.2 holds for some subsets of $\bar{U}_{i}$, while Condition 6.3 holds for some subsets of $\bar{U}_{i-1}$. Such cases are addressed by Corollaries 4.1 and 4.2 of Robins (1997); we give examples in $\S 7.2 .3$ below.

## 7. Influence diagrams

As previously mentioned, it is often helpful (though never essential) to represent and manipulate conditional independence properties graphically, using the formalism of influence diagrams (IDs). In particular, when including unobserved variables $\mathcal{U}$ and assuming extended stability, we can often deduce directly from graph-theoretic separation properties whether simple stability holds.

### 7.1. Semantics

Here we very briefly describe the semantics of IDs, and show how they can facilitate logical arguments by displaying implied properties in a particularly transparent form. We shall use the theory and notation of Cowell et al. (1999) and Dawid (2002) in relation to directed acyclic graphs (DAGs) and IDs, and their application to probability and decision models. The reader is referred to these sources for more details.

For any DAG or ID $\mathcal{D}$, its moral graph, or moralization, is the undirected graph $\operatorname{mo}(\mathcal{D})$ in which first an edge is inserted between any unlinked parents of a common child in $\mathcal{D}$, and then all directions are ignored. For any set $S$ of nodes of $\mathcal{D}$ we denote the smallest ancestral subgraph of $\mathcal{D}$ containing $S$ by $\operatorname{an}_{\mathcal{D}}(S)$, and its moralization by $\operatorname{man}_{\mathcal{D}}(S)$ (we may omit the specification of $\mathcal{D}$ when this is clear). For sets $A, B, C$ of nodes of $\mathcal{D}$ we write $A \Perp_{\mathcal{D}} B \mid C$, and say $C$ separates $A$ from $B$ (with respect to $\mathcal{D}$ ) to mean that, in man $(A \cup$ $B \cup C)$, every path joining $A$ to $B$ intersects $C$. Let $\operatorname{nd}(V)$ and pa $(V)$ denote the non-descendants and parents of a random node $V$, then it can be shown (Dawid, 2002; Lauritzen et al., 1990) that, whenever a probability distribution or decision problem is represented by $\mathcal{D}$, in the sense that for any such $V$ the probabilistic conditional independence $V \Perp \mathrm{nd}(V) \mid \mathrm{pa}(V)$ holds, we have

$$
\begin{equation*}
A \Perp_{\mathcal{D}} B|C \Rightarrow A \Perp B| C . \tag{31}
\end{equation*}
$$

This moralization criterion thus allows us to infer probabilistic independence properties from purely graph-theoretic separation properties. ${ }^{5}$

While the above allows us to read off conditional independencies from a DAG, we can, conversely, construct an ID $\mathcal{D}$ from a given collection of joint distributions over the domain variables (one for each regime) in the following way.

The node-set is given by $\mathcal{V}=\{\sigma\} \cup \mathcal{L} \cup \mathcal{U} \cup \mathcal{A}$. The graph has random (round) nodes for all the domain variables, and a founder decision (square) node for $\sigma$. The ordering given by the extended information base $\mathcal{I}^{\prime}$ induces an ordering on $\mathcal{V}$ such that any nodes in the (possibly empty) sets $L_{i}, U_{i}$ come after $A_{i-1}$ and before $A_{i}$, and $L_{N+1} \equiv Y$ is last. In addition we require the node $\sigma$ to be prior to any domain variables in this ordering. With each node $\nu \in \mathcal{V}_{0}:=\mathcal{V} \backslash\{\sigma\}$

[^4]is associated its collection of conditional distributions, given values for all its predecessors, pre $(\nu)$, in the ordering (including, in particular, specification of the relevant regime).

For each such $\nu$ we will have a conditional independence (CI) property of the form:

$$
C(\nu): \nu \Perp \operatorname{pre}(\nu) \mid \operatorname{pa}(\nu)
$$

where $\mathrm{pa}(\nu)$ is some given subset of pre $(\nu)$. Thus $C(\nu)$ asserts that the distributions of $\nu$, given all its predecessors, in fact only depends on the values of those in $\mathrm{pa}(\nu)$. Note that property $C(\nu)$ will be vacuous, and can be omitted, when $\operatorname{pa}(\nu)=\operatorname{pre}(\nu)$. Such a collection, $\mathcal{C}$ say, of CI properties is termed recursive. We represent $\mathcal{C}$ graphically by drawing an arrow into each node $\nu \in \mathcal{V}_{0}$ from each member of its parent set $\mathrm{pa}(\nu)$, and we associate with $\nu$ the 'parent-child' conditional probabilities of the form $p\left(\nu=\nu^{*} \mid \mathrm{pa}(\nu)=p a^{*}\right)$. The ID constructed in this way will ensure that the joint distribution of the domain variables, in each regime, satisfies any conditional independencies obtained by applying the moralization criterion (31).

From this point on, when we use the terms 'parents', 'ancestors' etc., the regime node $\sigma$ will be excluded from these sets. Also, while in general the terms $L_{i}, U_{i}$ could each refer to a collection of variables, for simplicity we shall consider only the case in which they represent just one (or sometimes none), and so can be modelled (if present at all) by a single node in the graph.

We emphasise that IDs are related to but distinct from 'causal DAGs' (Pearl, 1995; Spirtes, Glymour and Scheines, 2000). For a discussion see Dawid (2010) and Didelez, Kreiner and Keiding (2010).

### 7.2. Extended stability

The extended stability property (13) embodies a recursive collection of CI properties with respect to the ordering induced by the extended information base. Consequently it can be faithfully expressed by an ID $\mathcal{D}$ satisfying:

Condition 7.1 The only arrows out of $\sigma$ in $\mathcal{D}$ are into $\mathcal{A}$.
For $N=2$ this is depicted in Figure 2. Note that the subgraph corresponding to the domain variables is complete.

### 7.2.1. Sequential randomization

With the ordering induced by the extended information base $\mathcal{I}^{\prime}$, (13) and (15) together form a recursive collection $\mathcal{C}$ of CI properties. Therefore the conditions of Theorem 6.1 can be faithfully represented graphically in an ID $\mathcal{D}$, in which, for extended stability, the only arrows out of $\sigma$ are into the $A$ 's, while also, for sequential randomization, there are no arrows into the $A$ 's from the $U$ 's. Thus starting from Figure 2, for example, we simply delete all the arrows from a $U$ to an $A$, so obtaining Figure 3.


Fig 2. Unobserved variables: $N=2$.


Fig 3. ID showing sequential randomization.

We can now verify Theorem 6.1 using only graphical manipulations, as follows.

Since, under (13), the only children of $\sigma$ are action variables, and under (15) no action variable can be a child of any unobservable variable, it follows that in $\operatorname{man}\left(\sigma, \bar{L}_{i}, \bar{A}_{i-1}\right)$ there will be no direct link between $\sigma$ and any $U \in$ $\mathcal{U}$. A similar argument shows that (13) implies that there is no direct link in $\operatorname{man}\left(\sigma, \bar{L}_{i}, \bar{A}_{i-1}\right)$ between $\sigma$ and $L_{i}$. It follows that every path from $L_{i}$ to $\sigma$ must pass through one of the remaining variables, i.e. $\left(\bar{L}_{i-1}, \bar{A}_{i-1}\right)$, demonstrating that $L_{i} \Perp_{\mathcal{D}} \sigma \mid\left(\bar{L}_{i-1}, \bar{A}_{i-1}\right)$ for $i=1, \ldots, N+1$. Simple stability (6) now follows from (31).

### 7.2.2. Sequential irrelevance

The case of sequential irrelevance is more subtle. This is because when we combine extended stability (13) with sequential irrelevance (19) we do not obtain a recursive collection of CI properties. Consequently this combined collection of conditional independencies cannot be faithfully represented by any ID.

It might be thought that, starting with an ID representing extended stability, we could operate on it to incorporate (19) also simply by deleting all arrows from


FIG 4. ID implying sequential irrelevance.


Fig 5. Specialisations of Figure 4.
$U_{i}$ into $L_{j}$ for $j>i$. Doing this to Figure 2 yields the ID of Figure 4. However, that ID also represents the stronger property (30) (shown by the absence of edges from $\sigma$ and $\bar{U}_{i-1}$ into $L_{i}$ ), which does not follow from (13) and (19) without imposing further, non-graphical conditions (as was done in Theorem 6.2). We can indeed read off the stability property (6) from Figure 4, but while that graph thus displays clearly the conclusion of Theorem 6.2, it does not supply an alternative graphical proof.

By omitting some of the nodes and/or arrows in an ID, such as Figure 3 or Figure 4, that already embodies either sequential randomization or sequential irrelevance, we obtain simpler special cases with the same property. Two such examples, starting from Figure 4, are given in Figure 5.

### 7.2.3. Further examples

As mentioned before, we can have simple stability even when both sequential randomization and sequential irrelevance (or more precisely, the conditions of Theorems 6.1 and 6.2) fail. Two examples are given by the IDs of Figure 6. Applying the moralisation criterion to the graphs, we verify, for example, that in both IDs of Figure 6 simple stability is satisfied.

In full generality it is easy to see, using Condition 7.1, that application of the moralization criterion to $\mathcal{D}$ to check the simple stability condition (6) is


Fig 6. Alternative IDs displaying stability.
equivalent to checking that, for each $i, \bar{L}_{i-1}$ satisfies Pearl's back-door criterion (Pearl, 1995) relative to $\left(\bar{A}_{i-1}, L_{i}\right)$. (Pearl only considers atomic interventions, but our analysis shows that this condition also allows identification of conditional interventions.)

### 7.2.4. Positivity

Suppose that (whether by appealing to sequential randomization, or to sequential irrelevance, or the back-door criterion, or otherwise) we have been able to demonstrate simple stability with respect to an observable information base. Suppose further that $e$ is a control strategy in the sense of Condition 6.1. It will now follow that we can use $G$-recursion, exactly as in $\S 5.3$, to identify the consequence of regime $e$ from data gathered under regime $o$ - so long only as we can also ensure the positivity constraint of Definition 5.2.

It is easy to see that a sufficient condition for Definition 5.2 to hold is:
Condition 7.2 (Parent-child positivity) For each $A \in \mathcal{A}$, and each configuration $\left(a, p a^{*}\right)$ of $\left(A, \operatorname{pa}_{\mathcal{D}}(A)\right), p\left(a \mid p a^{*} ; e\right)>0 \Rightarrow p\left(a \mid p a^{*} ; o\right)>0$.

More generally, suppose that we specify, for each entry in each parent-child conditional probability table for the ID $\mathcal{D}$, whether it is zero or non-zero. We can then apply constraint propagation algorithms (Dechter, 2003) to determine $\mathcal{E}$ and $\mathcal{O}$. One such method (Dawid, 1992) uses an analogue of the computational method of probability propagation (Cowell et al., 1999). This generates a collection of 'cliques' (subsets of the variables) with, for each clique, an assignment of 1 (meaning possible) or 0 (impossible) to each configuration of its variables. Definition 5.2 will then hold if and only if, for each clique containing $\sigma$, no entry changes from 0 to 1 when we change the value of $\sigma$ from $o$ to $e$.

## 8. A more general approach

The simple stability condition (6) requires that, for each $i$, the conditional distribution of $L_{i}$, given the earlier variables $\left(\bar{L}_{i-1}, \bar{A}_{i-1}\right)$, should be the same under both regimes $o$ and $e$ - a strong assumption that, in certain problems, one
might be unwilling to accept directly, and unable to deduce, as in $\S 6.2$, from more acceptable assumptions. However, while we have shown that stability (together with Definition 5.2) is sufficient to support $G$-recursion, it turns out not to be necessary.

In this section we first give some very general conditions under which $G$ recursion can be justified; then we consider their specific application to models incorporating extended stability. Our analysis parallels parts of Robins (1987) (see also Section 3.4 of Robins (1997)), in which the 'sequential randomization' assumption is relaxed. We consider the relation between the two approaches in more detail in §10.2.

Rather than work directly with (10) and (11), we combine them into the following form:

$$
\begin{equation*}
f\left(\bar{l}_{i-1}, \bar{a}_{i-1}\right)=\sum_{l_{i}} \sum_{a_{i}} p\left(l_{i} \mid \bar{l}_{i-1}, \bar{a}_{i-1} ; o\right) \times p\left(a_{i} \mid \bar{l}_{i}, \bar{a}_{i-1} ; e\right) \times f\left(\bar{l}_{i}, \bar{a}_{i}\right) \tag{32}
\end{equation*}
$$

To justify $G$-recursion it is enough to demonstrate the applicability of (32).

### 8.1. G-recursion: General conditions

A primitive building block of our model is the specification of the interventional conditional probabilities $p\left(a_{i} \mid \bar{l}_{i}, \bar{a}_{i-1} ; e\right)$. We suppose that this is welldefined (e.g. by deterministic functions or specified randomization) at least for all $\left(\bar{l}_{i}, \bar{a}_{i-1}\right) \in \mathcal{O}(1 \leq i \leq N)$, even if $\left(\bar{l}_{i}, \bar{a}_{i-1}\right) \notin \mathcal{E}$.

We introduce a function $\gamma: \mathcal{H} \rightarrow\{0,1\}$ defined by:

$$
\gamma(h):= \begin{cases}1 & \text { if } h \in \mathcal{O} \text { and } \prod_{j=1}^{i} p\left(a_{j} \mid \bar{l}_{j}, \bar{a}_{j-1} ; e\right)>0  \tag{33}\\ 0 & \text { otherwise }\end{cases}
$$

In (33), $i$ is the highest index of an action variable appearing in $h$, i.e. $h=\left(\bar{l}_{i}, \bar{a}_{i}\right)$ or $\left(\bar{l}_{i+1}, \bar{a}_{i}\right)$. Note that if $h$ is an initial segment of $h^{\prime}$, then $\gamma(h)=0 \Rightarrow \gamma\left(h^{\prime}\right)=0$. We define:

$$
\begin{equation*}
\Gamma:=\{h \in \mathcal{H}: \gamma(h)=1\} \tag{34}
\end{equation*}
$$

(so that, in particular, $\Gamma \subseteq \mathcal{O}$ ).
We now impose the following positivity condition in place of Definition 5.2:
Condition 8.1 For $1 \leq i \leq N$, if $\left(\bar{l}_{i}, \bar{a}_{i-1}\right)$ is in $\Gamma$ and $p\left(a_{i} \mid \bar{l}_{i}, \bar{a}_{i-1} ; e\right)>0$, then $\left(\bar{l}_{i}, \bar{a}_{i}\right)$ is in $\mathcal{O}$ (and thus in $\Gamma$ ).

This requires that, subsequent to any partial history $\left(\bar{l}_{i}, \bar{a}_{i-1}\right)$ in $\Gamma$, if some value of the next action variable can be generated by intervention, it can also arise observationally.

Our approach now involves the construction, if possible, of a sequence of joint distributions $p_{i}(\cdot)(i=0, \ldots, N)$ for all the variables in the problem, such that

$$
\begin{equation*}
p_{0}(y) \equiv p(y ; e) \tag{35}
\end{equation*}
$$

and certain further properties hold, as described below. For maximum applicability these are stated here in a very abstract and general form. Some concrete cases where we can specify suitable $\left(p_{i}\right)$ and verify that they have the requisite properties are treated in $\S 8.2$ and $\S 10.2$ below.

Let the class of partial histories $h \in \mathcal{H}$ having positive probability under $p_{i}$ be denoted by $\mathcal{B}_{i}$, and let $\Gamma_{i}:=\mathcal{B}_{i} \cap \Gamma$.

We require the following positivity property:

$$
\begin{equation*}
\left(\bar{l}_{i}, \bar{a}_{i}\right) \in \mathcal{B}_{i} \Leftrightarrow\left(\bar{l}_{i}, \bar{a}_{i}\right) \in \mathcal{O} . \tag{36}
\end{equation*}
$$

Since $\Gamma \subseteq \mathcal{O}$, from ' $\Leftarrow$ ' in (36) we readily deduce

$$
\begin{equation*}
\left(\bar{l}_{i}, \bar{a}_{i}\right) \in \Gamma \Leftrightarrow\left(\bar{l}_{i}, \bar{a}_{i}\right) \in \Gamma_{i} . \tag{37}
\end{equation*}
$$

More substantively we require:

$$
\begin{align*}
p_{i-1}\left(l_{i} \mid \bar{l}_{i-1}, \bar{a}_{i-1}\right) & =p\left(l_{i} \mid \bar{l}_{i-1}, \bar{a}_{i-1} ; o\right) \quad(i=1, \ldots, N+1)  \tag{38}\\
p_{i-1}\left(a_{i} \mid \bar{l}_{i}, \bar{a}_{i-1}\right) & =p\left(a_{i} \mid \bar{l}_{i}, \bar{a}_{i-1} ; e\right) \quad(i=1, \ldots, N+1)  \tag{39}\\
p_{i-1}\left(y \mid \bar{l}_{i}, \bar{a}_{i}\right) & =p_{i}\left(y \mid \bar{l}_{i}, \bar{a}_{i}\right) \quad(i=1, \ldots, N) \tag{40}
\end{align*}
$$

whenever, in each case, the conditioning partial history on the left-hand side is in $\Gamma_{i-1}$ (in which case the conditional probabilities on both sides are unambiguously defined).

Suppose now that such a collection of distributions $\left(p_{i}\right)$ can be found. Let $\mathcal{H}_{0}$ denote the set of all partial histories of the form $\left(\bar{l}_{i}, \bar{a}_{i}\right)$ for some $i$. We define a function $f: \mathcal{H}_{0} \rightarrow \Re$ by:

$$
\begin{equation*}
f(h):=\gamma(h) \times \mathrm{E}_{i}\{k(Y) \mid h\}, \tag{41}
\end{equation*}
$$

for $h=\left(\bar{l}_{i}, \bar{a}_{i}\right)$, where $\mathrm{E}_{i}$ denotes expectation under $p_{i}$. We note that $f$ is welldefined, since $\gamma(h) \neq 0 \Rightarrow h \in \mathcal{O}$, whence $h \in \mathcal{B}_{i}$ by (36).

For $h=\left(\bar{l}_{N}, \bar{a}_{N}\right)$, if $\gamma(h) \neq 0$ then by (37) $h \in \Gamma_{N}$, so that we can apply (38) for $i=N+1$ to see that:

$$
f\left(\bar{l}_{N}, \bar{a}_{N}\right)= \begin{cases}\mathrm{E}\left\{k(Y) \mid \bar{l}_{N}, \bar{a}_{N} ; o\right\} & \text { if }\left(\bar{l}_{N}, \bar{a}_{N}\right) \in \Gamma  \tag{42}\\ 0 & \text { otherwise }\end{cases}
$$

Also, by (35),

$$
\begin{equation*}
f(\emptyset)=p(y ; e) . \tag{43}
\end{equation*}
$$

Lemma 8.1 Under Condition 8.1 and properties (35)-(40), the G-recursion (32) holds for the interpretation (41).

Proof. If $\gamma\left(\bar{l}_{i-1}, \bar{a}_{i-1}\right)=0$ then both sides of (32) are 0 .
Otherwise $\left(\bar{l}_{i-1}, \bar{a}_{i-1}\right)$ is in $\Gamma$ and so, by (37), in $\Gamma_{i-1}$. We have:

$$
\begin{align*}
f\left(\bar{l}_{i-1}, \bar{a}_{i-1}\right)= & \mathrm{E}_{i-1}\left\{k(Y) \mid \bar{l}_{i-1}, \bar{a}_{i-1}\right\} \\
= & \sum_{l_{i}} \sum_{a_{i}} p_{i-1}\left(l_{i} \mid \bar{l}_{i-1}, \bar{a}_{i-1}\right) \times p_{i-1}\left(a_{i} \mid \bar{l}_{i}, \bar{a}_{i-1}\right) \\
& \times \mathrm{E}_{i-1}\left\{k(Y) \mid \bar{l}_{i}, \bar{a}_{i}\right\} \tag{44}
\end{align*}
$$

Denote the three terms on the right-hand side of (44) by $T_{l}, T_{a}, T_{y}$, respectively. By (38) $T_{l}=p\left(l_{i} \mid \bar{l}_{i-1}, \bar{a}_{i-1} ; o\right)$. We do not need to consider the other terms when $T_{l}$ is 0 . Otherwise, $\left(\bar{l}_{i}, \bar{a}_{i-1}\right)$ is in $\Gamma_{i-1}$. By (39), we now have $T_{a}=p\left(a_{i} \mid\right.$ $\left.\bar{l}_{i}, \bar{a}_{i-1} ; e\right)$. Again we do not have to worry about $T_{y}$ unless $T_{a}$ is non-zero. In that case $\left(\bar{l}_{i}, \bar{a}_{i}\right)$ is in $\mathcal{B}_{i-1}$ and also, by Condition 8.1, in $\Gamma$, hence in $\Gamma_{i-1}$. We can now use (40) to replace $T_{y}$ by $\mathrm{E}_{i}\left\{k(Y) \mid \bar{l}_{i}, \bar{a}_{i}\right\}=f\left(\bar{l}_{i}, \bar{a}_{i}\right)$, and the result follows.

Starting from (42), we can thus apply $G$-recursion as given by (32), or equivalently by (10) and (11), to compute $f(\emptyset)$ - which, by (43), is just the desired consequence of regime $e$. In this computation we only need consider partial histories in $\Gamma$. When $e$ is a deterministic strategy we recover the form (12) of $G$-recursion.

Note that, for histories of intermediate length, the function $f$ defined by (41) involves the constructed distributions $\left(p_{i}\right)$, which need not have any realworld interpretation. Note further that, in contrast to the case when stability applies, even when we can use the above construction to compute the marginal interventional distribution of the response variable $Y$, there is no guarantee that we can identify the full joint interventional distribution of $\left(\bar{L}_{N}, \bar{A}_{N}, Y\right)$. In particular, if the loss function depends on variables other than $Y$ we may not be able to estimate the expected loss of an interventional strategy on the basis of observational data.

### 8.2. Extended stability

We now specialize the general approach of $\S 8.1$ to problems exhibiting extended stability, as in (13). This can be regarded as extending the analysis of Pearl and Robins (1995) to handle dynamic regimes, as also considered by Robins (1997). ${ }^{6}$

We aim to identify a graphical counterpart to the conditions of $\S 8.1$, that would allow us to apply $G$-recursion to this extended information base so as to identify the effect of regime $e$ from observations made under $o$.

For the remainder of this section we consider a given information base $\mathcal{I}^{\prime}$ that induces an ordering of the nodes of the influence diagram $\mathcal{D}$; in $\S 9$ we consider the converse, i.e. how to find an ordering of the information base from a given influence diagram $\mathcal{D}$ such that the graphical check of $\S 8.2 .1$ succeeds.

We impose Condition 7.2. It is then easy to see that Condition 8.1 will hold (and in fact $\Gamma=\mathcal{E}$ ). We also impose Condition 6.1 on the control strategy $e$.

For each $i=0, \ldots, N$, we now construct an artificial joint distribution $p_{i}(\cdot)$ for all the domain variables as follows. The distribution $p_{i}$ factors according to the ID $\mathcal{D}^{\prime}=\mathcal{D}$ with the node $\sigma$ removed. The parent-child tables for any variable $V \in \mathcal{L} \cup \mathcal{U}$ are unchanged from the original ones for $\mathcal{D}$ (which do not involve $\sigma)$. That for any action variable $A_{j}$ for $j \leq i$ is the same as for $\mathcal{D}$, conditional on $\sigma=o$; while that for $A_{j}(j>i)$ is the same as for $\mathcal{D}$, conditional on $\sigma=e$.

[^5]With this definition, $p_{0}(\cdot) \equiv p(\cdot ; e)$, so that (35) holds. Properties (36), and (38) for $i \leq N$, hold because the joint distribution of all variables up to and including $L_{i}$ is the same under $p_{i-1}$ as under $p(\cdot ; o)$; for (38) with $i=N+1$, when $L_{N+1} \equiv Y$, we also use the fact that extended stability, i.e. Condition 7.1, implies that the distribution of $Y$ given all earlier domain variables is the same under both $e$ and $o$.

Finally (39) holds because, by construction, the parent-child distribution for $A_{i}$ has the same specification for $p_{i-1}(\cdot)$ as for $p(\cdot ; e)$ - and, by Condition 6.1, $\operatorname{pa}\left(A_{i}\right) \subseteq\left(\bar{L}_{i}, \bar{A}_{i-1}\right)$.

### 8.2.1. Graphical check

We have shown that, under Conditions 6.1 and 7.2 , properties (35)-(39) hold automatically for our above construction of $\left(p_{i}\right)$. However, whether or not (40) holds will depend on more specific conditional independence properties of the problem under study. We now describe a graphical method based on IDs for checking this property.

For each action node $A \in \mathcal{A}$ we identify two subsets, $\mathrm{pa}_{o}(A)$ and $\mathrm{pa}_{e}(A)$, of $\operatorname{pa}_{\mathcal{D}}(A)$, such that, when $\sigma=o$ [resp. $\left.e\right]$, the conditional distribution of $A$, given its domain parents, can be chosen to depend only on pa ${ }_{o}(A)[$ resp. pa $e(A)]$.

To ensure Condition 6.1, we suppose:
Condition $8.2 \mathrm{pa}_{e}(A) \subseteq \mathcal{L} \cup \mathcal{A}$.
In order to investigate (40) for a specific value of $i$, we now construct, for $0 \leq i \leq N+1$, a new ID $\mathcal{D}_{i}$ on $\mathcal{V}$, as follows. The only arrow out of $\sigma$ (again a founder node) is now into $A_{i}$. For $j<i$, the parent set of $A_{j}$ is pa ${ }_{o}\left(A_{j}\right)$ with conditional distributions determined as under $o$; for $j>i$ it is $\mathrm{pa}_{e}\left(A_{j}\right)$, with conditional distributions determined as under $e$; finally, for $A_{i}$ it is ( $\left.\mathrm{pa}\left(A_{i}\right) ; \sigma\right)$, with conditional distributions exactly as in $\mathcal{D}$. Any domain variable $V \in \mathcal{L} \cup$ $\mathcal{U}$ has the same parent set $\mathrm{pa}(V)$ (which will not include $\sigma$ ) and conditional distributions as in $\mathcal{D}$. We shall use $\operatorname{an}_{i}(\cdot)$ to denote a minimal ancestral set in $\mathcal{D}_{i}$, with similar usages of $\mathrm{nd}_{i}, \Perp_{i}$, etc.

It is easy to see that the joint density of all the domain variables in $\mathcal{D}_{0}=\mathcal{D}_{e}$ is $p_{0}=p_{e}$; in $\mathcal{D}_{N+1}=\mathcal{D}_{o}$ it is $p_{N+1}=p_{o} ;$ while in $\mathcal{D}_{i}$, given $\sigma=o$ it is $p_{i-1}$, and given $\sigma=e$ it is $p_{i}$. Thus (40) will certainly hold if

$$
\begin{equation*}
Y \Perp_{i} \sigma \mid\left(\bar{L}_{i}, \bar{A}_{i}\right) \tag{45}
\end{equation*}
$$

holds. We can easily check (45) by inspection of the graph $\mathcal{D}_{i}$. Note that $\mathcal{D}_{0}$ is similar to the 'manipulated' DAG of Spirtes, Glymour and Scheines (2000).

In summary we have shown the following:
Theorem 8.2 Under Conditions 7.2 and 8.2, if the graphical separation property (45) holds for each $i$, then we can compute the consequence of regime $e$ from data gathered under regime o by means of the G-recursion (32), starting with $f_{N}$ as in (42), and ending with $f_{0}=p(y ; e)$.

A variant of this approach is described in Robins (1997), and works as follows. Let $\mathcal{D}_{i}^{\prime}$ be obtained from $\mathcal{D}_{i}$ by omitting the node $\sigma$, and deleting all arrows out of $A_{i}$. Because moralization links in $\mathcal{D}_{i}$ involving $\sigma$ can only be to predecessors of $A_{i}$, it is not difficult to see there exists a path from $Y$ to $\sigma$ avoiding $\left(\bar{L}_{i}, \bar{A}_{i}\right)$ in $\operatorname{man}_{\mathcal{D}_{i}}\left(Y, \bar{L}_{i}, \bar{A}_{i}\right)$ if and only if there exists such a path from $Y$ to $\mathrm{pa}\left(A_{i}\right)$ in $\operatorname{man}_{\mathcal{D}_{i}^{\prime}}\left(Y, \bar{L}_{i}, \bar{A}_{i}\right)$. And the latter condition can in turn be seen to be equivalent to the existence, in that graph, of a path from $Y$ to $A_{i}$ avoiding $\left(\bar{L}_{i}, \bar{A}_{i-1}\right)$. Thus $Y \Perp_{i} \sigma \mid\left(\bar{L}_{i}, \bar{A}_{i}\right)$ if and only if $Y \Perp_{\mathcal{D}_{i}^{\prime}} A_{i} \mid\left(\bar{L}_{i}, \bar{A}_{i-1}\right)$. Hence we can prove (40) by demonstrating the latter property.

It is shown in Dawid and Didelez (2008) that, under certain further conditions - informally, that each intermediate variable has some influence on the response under the interventional regime - when the graphical method described above succeeds we can deduce that the problem in fact exhibits simple stability with respect to the observed information base.

### 8.3. Examples

### 8.3.1. Stability

We first show that the conditions of $\S 5.2$ are a special case of those of $\S 8.1$, by verifying that the construction of $\S 8.2 .1$ works for the case of simple stability, as represented by Figure 1. In this case the $\left(U_{i}\right)$ are absent, and, for each domain variable $V, \mathrm{pa}_{e}(V)=\mathrm{pa}_{o}(V)=\operatorname{pre}(V)$. Thus $\mathcal{D}_{i}$ consists of the complete directed graph on ( $L_{1}, A_{1}, \ldots, L_{N}, A_{N}, Y$ ), together with an additional regime node $\sigma$ and an arrow from $\sigma$ to $A_{i}$. Figure 7 shows these graphs for the case $N=2$, and Figure 8 the corresponding graphs $\mathcal{D}_{i}^{\prime}$.

Since, after moralization of $\mathcal{D}_{i}, \sigma$ has direct links only into ( $\bar{L}_{i}, \bar{A}_{i}$ ), any path in this moral graph joining $Y$ to $\sigma$ must intersect $\left(\bar{L}_{i}, \bar{A}_{i}\right)$, whence we deduce (40). Equivalently, there is no path in $\mathcal{D}_{i}^{\prime}$ from $Y$ to $A_{i}$ avoiding $\left(\bar{L}_{i}, \bar{A}_{i-1}\right)$.


FIG 7. Influence diagrams $\mathcal{D}_{1}, \mathcal{D}_{2}$ for stability $(N=2)$.


FIG 8. Influence diagrams $\mathcal{D}_{1}^{\prime}, \mathcal{D}_{2}^{\prime}$ for stability $(N=2)$.

Hence we have confirmed that, when stability holds, it is possible to construct a sequence of joint densities $p_{i}$ satisfying (38)-(40).

### 8.3.2. G-recursion without stability

More interesting is the possibility of applying the construction of $\S 8.2$ to justify $G$-recursion even in cases where simple stability does not hold. This is illustrated by the following example, based on Pearl and Robins (1995) (and see Robins (1987) and Robins (1997) for description of medical scenarios that are reasonably captured by this example).

Example 8.1 Figure 9 shows a specific model incorporating extended stability for the information base $\left(U_{1}, A_{1}, U_{2}, L_{2}, A_{2}, Y\right)$ (with $L_{1}=\emptyset$ ). Note that this does not embody simple stability, since moralization would create a direct link between $\sigma$ and $U_{1}$, and hence a path $L_{2}-U_{1}-\sigma$ that avoids $A_{1}$. We thus can not deduce $L_{2} \Perp \sigma \mid A_{1}$, as would be required for simple stability.

We use stippled arrows to represent independence under the control regime $e$. Thus the stippled arrow from $U_{1}$ to $A_{1}$ in Figure 9 represents the property

$$
\begin{equation*}
A_{1} \Perp U_{1} \mid \sigma=e, \tag{46}
\end{equation*}
$$

which is (14) for $i=1$. (The equivalent property for $i=2$ is already implied by the lack of any arrows from $U_{1}$ and $U_{2}$ to $A_{2}$ ).

The stippled arrow from $L_{2}$ to $A_{2}$ embodies an additionally assumed property:

$$
\begin{equation*}
A_{2} \Perp L_{2} \mid\left(A_{1} ; \sigma=e\right) . \tag{47}
\end{equation*}
$$

That is, we are supposing that interventional assignment of $A_{2}$ can only depend (deterministically or stochastically) on the value chosen for the previous treatment, $A_{1}$. This is a restriction on the type of interventional strategy $e$ that we are considering. It will turn out that we can identify the causal effect of $e$ from the observational data gathered under $o$, using $G$-recursion, only for strategies $e$ of this special type.

In this problem we thus have pa ${ }_{o}\left(A_{1}\right)=U_{1}, \mathrm{pa}_{e}\left(A_{1}\right)=\emptyset, \mathrm{pa}_{o}\left(A_{2}\right)=\left(A_{1}, L_{2}\right)$, $\mathrm{pa}_{e}\left(A_{2}\right)=A_{1}$. The constructed IDs $\mathcal{D}_{1}$ and $\mathcal{D}_{2}$ are shown in Figure 10, and the variant forms $\mathcal{D}_{1}^{\prime}$ and $\mathcal{D}_{2}^{\prime}$ (Pearl and Robins, 1995, Figure 2) in Figure 11.


Fig 9. An ID displaying non-stability.


FIG 10. Influence diagrams $\mathcal{D}_{1}, \mathcal{D}_{2}$ for Figure 9.


FIG 11. Influence diagrams $\mathcal{D}_{1}^{\prime}, \mathcal{D}_{2}^{\prime}$ for Figure 9.


FIG 12. Relevant moral ancestral graphs, for $\mathcal{D}_{1}$ and $\mathcal{D}_{1}^{\prime}$.
We first examine $\mathcal{D}_{1}$ to see if $Y \Perp_{\mathcal{D}_{1}} \sigma \mid A_{1}$. The relevant moral ancestral graph (see Figure 12) is easily seen to have the desired separation property: thus we have shown (40) for $i=1$. Alternatively, from examination of the relevant moral ancestral graph based on $\mathcal{D}_{1}^{\prime}$ we readily see the desired property $Y \Perp_{\mathcal{D}_{1}^{\prime}} A_{1}$. (Note that this approach does not succeed if we allow $A_{2}$ to depend on $L_{2}$ under $e$, thus retaining an arrow from $L_{2}$ to $A_{2}$ and so making $L_{2}$ an ancestor of $Y$ in $\mathcal{D}_{1}$ : in the now larger relevant moral ancestral graph formed from $\mathcal{D}_{1}$ we could then trace a path $Y-U_{2}-U_{1}-\sigma$ from $Y$ to $\sigma$ avoiding $A_{1}$.)

Finally, since in $\mathcal{D}_{2}$ neither $U_{1}$ nor $U_{2}$ is a parent of $A_{2}$, even after moralization there will be no direct link from $\sigma$ to either $U_{1}$ or $U_{2}$ : consequently any path from $Y$ to $\sigma$ will have to intersect $\left(A_{1}, L_{2}, A_{2}\right)$. Equivalently, we see that in $\mathcal{D}_{2}^{\prime}$, after moralization (which adds a futher link between $U_{1}$ and $U_{2}$ ) every path from $Y$ to $A_{2}$ intersects $\left(A_{1}, L_{2}\right)$. We deduce $Y \Perp \sigma \mid\left(A_{1}, L_{2}, A_{2}\right)$, i.e. (40) for $i=2$.

If we now assume Conditions 7.2 and 6.1 then, all the required conditions being satisfied, we will have justified use of $G$-recursion to identify the consequences of an interventional regime $e$ of the specified form, from data collected under the observational regime $o$.

The graphical check illustrated above simplifies considerably in the case of an unconditional interventional strategy $e$, where the values of the action variables are determined in advance, as considered by Pearl and Robins (1995). In this case $\operatorname{pa}_{e}\left(A_{i}\right)=\emptyset$ for all $i$, and $\mathcal{D}_{i}$ is obtained from $\mathcal{D}$ by deleting all arrows into every $A_{j}$ with $j>i$. Then $\mathcal{D}_{i}^{\prime}$ is obtained by further deleting $\sigma$ and all arrows out of $A_{i}$. However, if our aim is to compare strategies, and ideally find an optimal one, it is necessary also to consider dynamic strategies.

## 9. Constructing an admissible sequence

In order to apply the graphical check of $\S 8.2 .1$ we need to have the variables already completely ordered. More generally, we could ask whether there exists an ordering $\left(A_{1}, \ldots, A_{N}\right)$ of $\mathcal{A}$, and $\left(L_{1}, \ldots, L_{N}\right)$ of disjoint subsets of $\mathcal{L}$, such that we can apply the construction of $\S 8.2 .1$ to show (45). Somewhat more restricted, we might suppose an ordering $\left(A_{1}, \ldots, A_{N}\right)$ already given, and look for a sequence $\left(L_{1}, \ldots, L_{N}\right)$ to satisfy (45). Such a sequence will be termed admissible. In this section we assume that a graphical representation of the problem in form of an ID is given, and we note that by definition an admissible sequence has to satisfy $\bar{L}_{i} \subseteq \operatorname{nd}\left(A_{i}, \ldots, A_{N}\right)$. Below, we give conditions under which we can determine whether such an admissible sequence exists, and construct one if it does. We shall need some general properties of directed-graph separation from Appendix A.

We impose the following conditions:
Condition 9.1 For all $i$,

$$
\operatorname{pa}_{e}\left(A_{i}\right) \subseteq \operatorname{pa}_{o}\left(A_{i}\right)
$$

This can always be ensured by redefining, if necessary, $\mathrm{pa}_{o}\left(A_{i}\right)$ as $\mathrm{pa}_{o}\left(A_{i}\right) \cup$ $\mathrm{pa}_{e}\left(A_{i}\right)$, with any added parents having no effect on the conditional probabilities for $A_{i}$ under $o$.

Condition 9.2 Each action variable $A \in \mathcal{A}$ is an ancestor of $Y$ in $\mathcal{D}_{e}$.
In typical contexts Condition 9.2 will hold, since we would not normally contemplate an intervention that has no effect on the response. Clearly when Conditions 9.1 and 9.2 both hold every $A \in \mathcal{A}$ is also an ancestor of $Y$ in $\mathcal{D}_{o}=\mathcal{D}$.

Define, for $i=1, \ldots, N$ :

$$
\begin{equation*}
M_{i}:=\mathcal{L} \cap \operatorname{nd}_{e}\left(A_{i}, A_{i+1}, \ldots, A_{N}\right) \cap \operatorname{an}_{i}(Y) \tag{48}
\end{equation*}
$$

We note that $M_{i-1} \subseteq M_{i}$. This follows from $\operatorname{an}_{i-1}(Y) \subseteq \mathrm{an}_{i}(Y)$ which in turn holds because, by Condition 9.1, the edge set of $\mathcal{D}_{i-1}$ is a subset of that of $\mathcal{D}_{i}$.

Now let

$$
\begin{equation*}
L_{i}^{*}:=M_{i} \backslash M_{i-1} \tag{49}
\end{equation*}
$$

so that $M_{i}=\bar{L}_{i}^{*}$. For the information sequence $\left(L_{i}^{*}\right)$, the total information taken into account up to time $i, M_{i}$, consists of just those variables in $\mathcal{L}$ that are ancestors of $Y$ in $\mathcal{D}_{i}$, but are not descendants of $A_{i}$ or any later actions.

The sequence $\left(L_{1}^{*}, \ldots, L_{N}^{*}\right)$ will be admissible if, for $i=1, \ldots, N$,

$$
\begin{equation*}
Y \Perp_{i} \sigma \mid\left(M_{i}, \bar{A}_{i}\right) \tag{50}
\end{equation*}
$$

Taking into account Condition 9.2 and (48), (50) requires that $M_{i} \cup \bar{A}_{i}$ separate $Y$ from $\sigma$ in the undirected graph $\mathcal{G}_{i}$ obtained by moralizing the ancestral set of $Y$ in $\mathcal{D}_{i}$. It is thus straightforward to check whether or not it holds. When it does we shall call $i$ admissible.

The following result can be regarded as simultaneously simplifying, generalizing, and rendering more operational that of Pearl and Robins (1995). In particular, it supplies an explicit construction, while allowing for conditional interventions.

Theorem 9.1 Under Conditions 9.1 and 9.2, if any admissible sequence exists then $\left(L_{1}^{*}, \ldots, L_{N}^{*}\right)$ is admissible.

That is: There exists an admissible sequence if and only if every $i$ is admissible. In this case $\left(L_{1}^{*}, \ldots, L_{N}^{*}\right)$ is an admissible sequence.

Proof. Suppose that there exists some admissible sequence $\left(L_{1}, \ldots, L_{N}\right)$. Then, for each $i$,

$$
\begin{equation*}
Y \Perp_{i} \sigma \mid \bar{L}_{i} \cup \bar{A}_{i} . \tag{51}
\end{equation*}
$$

By Lemma A.2, this graph-theoretical separation continues to hold if we intersect the conditioning set with $\left(\operatorname{an}_{i}(Y), \sigma\right)$. Since, by Condition 9.2, $\bar{A}_{i} \subseteq \operatorname{an}_{i}(Y)$, we obtain

$$
\begin{equation*}
Y \Perp_{i} \sigma \mid\left(\bar{L}_{i} \cap \operatorname{an}_{i}(Y)\right) \cup \bar{A}_{i} . \tag{52}
\end{equation*}
$$

But, $\bar{L}_{i} \subseteq \mathcal{L} \cap \operatorname{nd}_{e}\left(A_{i}, \ldots, A_{N}\right) ;$ and thus $\bar{L}_{i} \cap \mathrm{an}_{i}(Y) \subseteq M_{i}$. Hence, by Lemma A.2, (50) holds, and the result follows.

Example 9.1 (We are indebted to Susan Murphy for this example.) In the problem represented in Figure 13, it may be checked that the 'obvious' choice $L_{1}=\{X\}, L_{2}=\{Z\}$ is not an admissible sequence. Using the method above


Fig 13. Finding an admissible sequence.
yields $L_{1}^{*}=\{X, Z\}, L_{2}^{*}=\emptyset$, which is admissible (indeed, yields simple stability, as may either be checked directly, or deduced from Theorem 2 in Dawid and Didelez (2008)).

### 9.1. Finding a better sequence

While the above procedure will always construct an admissible sequence ( $L_{1}, \ldots$, $\left.L_{N}\right)$ when one exists, that might not be the best possible. Thus in Figure 14, with $\mathcal{L}=\{X, Z\}$, we find $L_{1}^{*}=\{Z\}, L_{2}^{*}=\{X\}$. These satisfy (50), so that the sequence $\left\{L_{1}^{*}, L_{2}^{*}\right\}$ is admissible. However a smaller admissible sequence is given by $L_{1}=\emptyset, L_{2}=\{X\}$.

If we had initially regarded $Z$ as unobservable, so taking $\mathcal{L}=\{X\}$, we would have found this smaller sequence. However in general we would need hindsight or good fortune to start off with such a minimal specification of $\mathcal{L}$.

Even without redefining $\mathcal{L}$, however, we can often improve on the sequence given by (49). At each stage $i$ we first check (50). If this fails we abort the process. Otherwise, sequentially choose $L_{i}$ to be any subset of $M_{i}$, disjoint from $\bar{L}_{i-1}$, such that (51) holds. (Since, by (50), (51) holds for the choice $L_{i}=M_{i} \backslash \bar{L}_{i-1}$, such a set must exist.) Then (if the process is never aborted) we shall have constructed an admissible sequence $\left(L_{i}\right)$, improving on $\left(L_{i}^{*}\right)$ in the sense that $\bar{L}_{i} \subseteq \bar{L}_{i}^{*}$.

Ideally we would want the set $L_{i}$ to be small. When each $L_{i}$ is minimal, in the sense that no proper subset of $L_{i}$ satisfies (51), we obtain a generalization of the method of Pearl and Robins (1995) for constructing a minimal admissible sequence. However in large problems the search for such a minimal $L_{i}$ can be computationally non-trivial, and we may have to be satisfied with some other choices for the $\left(L_{i}\right)$. Minimality is in any case not a requirement for admissibility.

### 9.2. Admissible orderings of $\mathcal{A}$

In general there will be several orderings of $\mathcal{A}$ possible. It can then happen that an admissible sequence $\left(L_{1}, \ldots, L_{N}\right)$ exists for one ordering of $\mathcal{A}$ (which we may then likewise call admissible), but not for another.


Fig 14. A choice of admissible sequences.


FIG 15. Unordered actions.
Example 9.2 In the ID of Figure $15, \mathcal{U}=\{U\}, \mathcal{L}=\{L\}, \mathcal{A}=\{A, B\}$. Note that $A \Perp B$ under either regime. Both $A_{1}=A, A_{2}=B$ and $A_{1}=B, A_{2}=A$ are possible orderings of $\mathcal{A}$. For the former choice we find $M_{1}=\emptyset$; then (50) for $i=1$ becomes $Y \Perp_{\mathcal{D}_{A}} \sigma \mid A$, where $\mathcal{D}_{A}$ is $\mathcal{D}$ with the arrow from $\sigma$ to $B$ removed. Since this is easily seen to fail (moralization creates a link between $U$ and $\sigma$ ), Theorem 9.1 implies that there can be no admissible sequence to support $G$-recursion. However if we take $A_{1}=B, A_{2}=A$, we obtain $M_{1}=\emptyset$, $M_{2}=\{L\}$, and (50) becomes $Y \Perp_{\mathcal{D}_{B}} \sigma \mid B$ for $i=1, Y \Perp_{\mathcal{D}_{A}} \sigma \mid(B, L, A)$ for $i=2$, where $\mathcal{D}_{B}$ is $\mathcal{D}$ with the arrows into $A$ from both $\sigma$ and $U$ removed. Both of these properties are easily confirmed to hold. We can thus (under suitable positivity conditions) apply $G$-recursion with respect to the admissible ordering ( $B, L, A$ ).

As yet we do not have a method that will automatically identify an admissible ordering of $\mathcal{A}$ when one exists.

## 10. Potential response models

In this section, we examine the relationship between the potential response ( $P R$ ) approach to dynamic treatments and our own decision-theoretic one.

The PR approach typically confines attention to non-randomized, though possibly dynamic, strategies. Such a strategy is defined by a function $g$ on the set of all 'partial $L$-histories' of the form $\left(\bar{l}_{i}\right)(1 \leq i \leq N)$, such that, for each $i$, $g\left(\bar{l}_{i}\right)$ is one of the available options for $A_{i}$. We shall write $\bar{g}\left(\bar{l}_{i}\right)$ for the sequence $\left(g\left(l_{1}\right), g\left(l_{1}, l_{2}\right), \ldots, g\left(\bar{l}_{i}\right)\right)$. Under this strategy, if at time $i$ we have observed $\bar{L}_{i}=\bar{l}_{i}$, the next action will be $A_{i}=g\left(\bar{l}_{i}\right)$.

We henceforth confine attention to a pair of regimes $\mathcal{S}=\{o, e\}$, where $o$ is observational, while $e$ is a non-randomized strategy, determined by a given function $g$ as described above.

### 10.1. Potential responses and stability

We first interpret and analyse the model introduced by Robins (1986) (see also Robins (1997), Section 3.3; Robins (2000); Murphy (2003)).

We need to introduce, for each regime $s \in\{o, e\}$, a collection of 'potential variables' $\Pi_{s}:=\left(L_{s, 1}, A_{s, 1}, \ldots, L_{s, N}, A_{s, N}, L_{s, N+1} \equiv Y_{s}\right)$. It is supposed that, when regime $s$ is operating, the actual observable variables in the problem, $\left(L_{1}, A_{1}, \ldots, L_{N}, A_{N}, Y\right)$, will be those in $\Pi_{s}$.

Note that, by the definition of $e$, we have the functional constraint

$$
\begin{equation*}
A_{e, i}=g\left(\bar{L}_{e, i}\right) \quad(i=1, \ldots, N) \tag{53}
\end{equation*}
$$

All the potential variables, across both regimes, are regarded as having simultaneous existence, their values being unaffected by which regime is actually followed. ${ }^{7}$ The effect of following regime $s$ is thus to uncover the values of some of these, viz. those in $\Pi_{s}$, while hiding others.

This collection of all potential observables across both regimes is further considered to have a joint distribution (respecting the logical constraints (53)), whose density we denote by $p(\cdot)$. This distribution is supposed unaffected by which regime is in operation: all this can do is change the relationship between potential and actual variables.

Since, under $e, Y \equiv Y_{e}$, the consequence of the interventional strategy $e$ is simply the marginal distribution of $Y_{e}$. Our aim is to identify this distribution from observations made under regime $o$.

It can be shown directly that this can be effected by means of the $G$-recursion formula under the following conditions:
Condition 10.1 (Positivity) Whenever $p\left(\bar{L}_{o, N}=\bar{l}_{N}\right)>0$,

$$
p\left(\bar{A}_{o, N}=\bar{g}\left(\bar{l}_{N}\right) \mid \bar{L}_{o, N}=\bar{l}_{N}\right)>0
$$

That is, in the observational regime, for any set of values $\bar{l}_{N}$ of the variables $\bar{L}_{N}$ that can arise with positive probability, there is a positive probability that the actions taken will be those specified by $e$.
Condition 10.2 (Consistency) If $\bar{A}_{o, i}=\bar{g}\left(\bar{L}_{o, i}\right)$, then $L_{o, i+1}=L_{e, i+1}(i=$ $0, \ldots, N)$.
(Note that for $i=0$ the antecedent of this condition is vacuously satisfied, while for $i=N$ its conclusion is $Y_{o}=Y_{e}$.)

That is, if, in the observational regime, we happen to obtain a partial history $\left(\bar{l}_{i}, \bar{a}_{i}\right)$ that could also be obtained under the operation of $e$, then we will next observe the identical variable $L_{e, i+1}$ that would have been observed if we had been operating $e$. (This condition of course imposes further logical constraints on the joint distribution $p$ ).
Condition 10.3 (Sequential ignorability) Whenever $p\left(\bar{L}_{o, i}=\bar{l}_{i}\right)>0$,

$$
A_{o, i} \Perp \bar{L}_{e}^{i+1} \mid\left(\bar{L}_{o, i}=\bar{l}_{i}, \bar{A}_{o, i-1}=\bar{g}\left(\bar{l}_{i-1}\right)\right) \quad(i=1, \ldots, N),
$$

$\left(\right.$ where $\bar{L}_{e}^{j}:=\left(L_{e, j}, L_{e, j+1}, \ldots, L_{e, N}, Y_{e}\right)$ ).

[^6]That is, in the observational regime, given any partial history consistent with the operation of $e$, the next action is independent of all the future potential observables associated with $e .{ }^{8}$

### 10.1.1. Connexions

We now consider the relationship between the above approach and that of §5.2, which founds $G$-recursion on the stability property (6). We will show that Conditions $10.1,10.2$ and 10.3 imply our conditions in $\S 5.2$. Our reasoning is, in spirit, very similar to Theorem 3.1 of Robins (1997) (see also Robins (1986), Theorem 4.1).
Lemma 10.1 If Conditions 10.2 and 10.3 hold, then for any sequence $\bar{l}_{N+1}=$ $\left(l_{1}, \ldots, l_{N}, y\right)$ such that $p\left(\bar{L}_{e, N}=\bar{l}_{N}\right)>0$,

$$
\begin{equation*}
p\left(\bar{L}_{e}^{i+1}=\bar{l}^{i+1} \mid \bar{L}_{e, i}=\bar{l}_{i}, \bar{A}_{o, i}=\bar{g}\left(\bar{l}_{i}\right)\right)=p\left(\bar{L}_{e}^{i+1}=\bar{l}^{i+1} \mid \bar{L}_{e, i}=\bar{l}_{i}\right) \tag{54}
\end{equation*}
$$

for $i=0, \ldots, N$.
Proof. First note that, from Condition 10.2, when $\bar{A}_{o, i}=\bar{g}\left(\bar{l}_{i}\right), \bar{L}_{o, i+1}=\bar{l}_{i+1}$ is equivalent to $\bar{L}_{e, i+1}=\bar{l}_{i+1}$. So from Condition 10.3

$$
\begin{equation*}
A_{o, i+1} \Perp \bar{L}_{e}^{i+2} \mid\left(\bar{L}_{e, i+1}=\bar{l}_{i+1}, \bar{A}_{o, i}=\bar{g}\left(\bar{l}_{i}\right)\right) . \tag{55}
\end{equation*}
$$

We now show (54) by induction on $i$.
It holds trivially for $i=0$. Suppose then it holds for $i$. Conditioning both sides on $L_{e, i+1}=l_{i+1}$ then yields

$$
p\left(\bar{L}_{e}^{i+2}=\bar{l}^{i+2} \mid \bar{L}_{e, i+1}=\bar{l}_{i+1}, \bar{A}_{o, i}=\bar{g}\left(\bar{l}_{i}\right)\right)=p\left(\bar{L}_{e}^{i+2}=\bar{l}^{i+2} \mid \bar{L}_{e, i+1}=\bar{l}_{i+1}\right)
$$

But from (55) we have

$$
\begin{aligned}
& p\left(\bar{L}_{e}^{i+2}=\bar{l}^{i+2} \mid \bar{L}_{e, i+1}=\bar{l}_{i+1}, \bar{A}_{o, i+1}=\bar{g}\left(\bar{l}_{i+1}\right)\right) \\
& \quad=p\left(\bar{L}_{e}^{i+2}=\bar{l}^{i+2} \mid \bar{L}_{e, i+1}=\bar{l}_{i+1}, \bar{A}_{o, i}=\bar{g}\left(\bar{l}_{i}\right)\right)
\end{aligned}
$$

Hence (54) holds with $i$ replaced by $i+1$ and the induction proceeds.

Theorem 10.2 If Conditions 10.2 and 10.3 hold, then so does the stability condition (6).
Proof. Because of (53), and the restriction immediately below the density interpretation (7) of (6), it is enough to show that $p\left(L_{e, i+1}=l_{i+1} \mid \bar{L}_{e, i}=\right.$ $\left.\bar{l}_{i}, \bar{A}_{e, i}=\bar{g}\left(\bar{l}_{i}\right)\right)=p\left(L_{o, i+1}=l_{i+1} \mid \bar{L}_{o, i}=\bar{l}_{i}, \bar{A}_{o, i}=\bar{g}\left(\bar{l}_{i}\right)\right)$. But, again by (53), $p\left(L_{e, i+1}=l_{i+1} \mid \bar{L}_{e, i}=\bar{l}_{i}, \bar{A}_{e, i}=\bar{g}\left(\bar{l}_{i}\right)\right)=p\left(L_{e, i+1}=l_{i+1} \mid \bar{L}_{e, i}=\bar{l}_{i}\right)$. By

[^7]Lemma 10.1, this is the same as $p\left(L_{e, i+1}=l_{i+1} \mid \bar{L}_{e, i}=\bar{l}_{i}, \bar{A}_{o, i}=\bar{g}\left(\bar{l}_{i}\right)\right)$, and by Condition 10.2 this is in turn the same as $p\left(L_{o, i+1}=l_{i+1} \mid \bar{L}_{e, i}=\bar{l}_{i}, \bar{A}_{o, i}=\right.$ $\left.\bar{g}\left(\bar{l}_{i}\right)\right)$.

Finally, in the light of (53), it is easy to see that Condition 10.1 implies positivity as given by Definition 5.2.

In summary, whenever the conditions usually used to justify $G$-recursion in the potential response framework hold, so will our own (as in §5.2). But our conditions are more general in that they do not require the existence of, let alone any probabilistic relationships between, potential responses under different regimes; and can, moreover, just as easily handle randomized interventional strategies, which are more problematic for the PR approach.

### 10.2. Potential responses without stability

A more general approach (Gill and Robins, 2001; Lok et al., 2004; Robins, 1987, 1989; Robins, Hernán and Brumback, 2000) within the potential response framework replaces Conditions 10.2 and 10.3 with the following variants:
Condition 10.4 If $\bar{A}_{o, N}=\bar{g}\left(\bar{L}_{o, N}\right)$, then $Y_{o}=Y_{e}$.
That is, if in the observational regime we happen to observe a complete history that could have arisen under the operation of $e$, then the response will be identical to what we would have observed had we been operating $e$.

## Condition 10.5

$$
A_{o, i} \Perp Y_{e} \mid\left(\bar{L}_{o, i}=\bar{l}_{i}, \bar{A}_{o, i-1}=\bar{g}\left(\bar{l}_{i-1}\right)\right) \quad(i=1, \ldots, N) .
$$

That is, if, in the observational strategy, we happen to observe a partial history that could have arisen under the operation of $e$, then the next action is independent of the potential response under $e$.

Condition 10.4 implies, and can in fact be replaced by:
Condition 10.6 Given $\left(\bar{L}_{o, N}=\bar{l}_{N}, \bar{A}_{o, N}=\bar{g}\left(\bar{l}_{N}\right)\right), Y_{o}$ and $Y_{e}$ have the same conditional distribution.

The deterministic strategy $e$ is termed evaluable if, for each $i$ :
Condition 10.7

$$
p\left(\bar{L}_{o, i}=\bar{l}_{i}, \bar{A}_{o, i-1}=g\left(\bar{l}_{i-1}\right)\right)>0 \Rightarrow p\left(\bar{L}_{o, i}=\bar{l}_{i}, \bar{A}_{o, i}=g\left(\bar{l}_{i}\right)\right)>0 .
$$

Note that Conditions 10.4-10.7 make no mention of potential intermediate variables ( $\bar{L}_{e, N}, \bar{A}_{e, N}$ ) under $e$ - though they do involve both versions, $Y_{o}$ and $Y_{e}$, of the response. The relevant variables in the problem can thus be taken as ( $\bar{L}_{o, N}, \bar{A}_{o, N}, Y_{o}, Y_{e}$ ), having a joint distribution $p$ say.

Conditions 10.5 and 10.6 are weaker than those of $\S 10.1$ as none of the variables under strategy $e$ other than $Y_{e}$ are involved. Note that, for example,
it is not required that, when an observational partial history could have arisen under $e$, that is the history that would have so arisen; but even so, constraints on $Y_{e}$ are then imposed.

### 10.2.1. Connexions

It is straightforward to show directly that, when Conditions 10.5, 10.6 and 10.7 hold, the marginal distribution of $Y_{e}$, or the interventional consequence $\mathrm{E}\left\{k\left(Y_{e}\right)\right\}$, can be identified by the $G$-recursion (12). We now show how this approach can be related to our own decision-theoretic one. Specifically, we shall show that, when the above conditions hold, so do those of $\S 8.1$ (see also Theorem 3.2 of Robins (1997)).

Condition 10.7 is just Condition 8.1 specialized to the case of the deterministic strategy $e$.

To continue, we construct a fictitious distribution $p_{i}(\cdot)(i=0, \ldots, N)$, for variables $\left(L_{1}, A_{1}, \ldots, L_{N}, A_{N}, Y\right)$, as follows.

Definition 10.1 The distribution $p_{i}$ of ( $L_{1}, A_{1}, \ldots, L_{N}, A_{N}, Y$ ) is defined as the distribution under $p$ of $\left(L_{o, 1}, A_{o, 1}, \ldots, L_{o, i}, A_{o, i}, L_{o, i+1}, g\left(\bar{L}_{o, i+1}\right), \ldots, L_{o, N}\right.$, $\left.g\left(\bar{L}_{o, N}\right), Y_{e}\right)$.
Thus

$$
\begin{align*}
& p_{i}\left(\bar{L}_{N}=\bar{l}_{N}, \bar{A}_{N}=\bar{a}_{N}, Y=y\right):= \\
& \qquad \begin{cases}p\left(\bar{L}_{o, N}=\bar{l}_{N}, \bar{A}_{o, i}=\bar{a}_{i}, Y_{e}=y\right) & \text { if } a_{i+1}=g\left(\bar{l}_{i+1}\right), \ldots, a_{N}=g\left(\bar{l}_{N}\right) \\
0 & \text { otherwise }\end{cases} \tag{56}
\end{align*}
$$

Note that this construction of $p_{i}$ is quite different from that developed, in a different context, in $\S 8.2$. In particular, the marginal joint distribution of $\left(\bar{L}_{N}\right)$ is, for every $p_{i}$, the same as under $p_{o}$.

Equation (35) follows trivially from Definition 10.1.
As in $\S 8.2$, Properties (36), and (38) for $i \leq N$, hold because the joint distribution of all variables up to and including $L_{i}$ is the same under $p_{i-1}$ as under $p(\cdot ; o)$; while for (38) with $i=N+1$, when $L_{N+1} \equiv Y$, we also use Condition 10.6.

Equation (39) holds since the distribution on either side is concentrated on $g\left(\bar{l}_{i}\right)$.

Finally we show (40).
We only need this for $\left(\bar{l}_{i}, \bar{a}_{i}\right) \in \Gamma_{i-1}$. Since then $\left(\bar{l}_{i}, \bar{a}_{i}\right) \in \Gamma$, we must by (33) have $p\left(a_{j} \mid \bar{l}_{j}, \bar{a}_{j-1} ; e\right)>0(1 \leq j \leq i)$, which in virtue of the deterministic nature of strategy $e$ requires $\bar{a}_{i}=\bar{g}\left(\bar{l}_{i}\right)$; and then the additional condition $\left(\bar{l}_{i}, \bar{a}_{i}\right) \in \mathcal{O}$ becomes $p\left(\bar{L}_{o, i}=\bar{l}_{i}, \bar{A}_{o, i}=\bar{g}\left(\bar{l}_{i}\right)\right)>0$. So in this case (40) becomes:

$$
\begin{equation*}
p\left(Y_{e}=y \mid \bar{L}_{o, i}=\bar{l}_{i}, \bar{A}_{o, i-1}=\bar{g}\left(\bar{l}_{i-1}\right)\right)=p\left(Y_{e}=y \mid \bar{L}_{o, i}=\bar{l}_{i}, \bar{A}_{o, i}=\bar{g}\left(\bar{l}_{i}\right)\right) \tag{57}
\end{equation*}
$$

But this is an immediate consequence of Condition 10.5.
In summary, we have shown:

Theorem 10.3 Under Conditions 10.1, 10.6 and 10.5, and defining $p_{i}(\cdot)$ by Definition 10.1, Conditions 5.2 and 8.1 and equations (35)-(40) are all satisfied.

From Lemma 8.1 we now deduce:
Corollary 10.4 Under Conditions 10.5-10.7, the consequence of strategy e can be recovered using the $G$-recursion (12).

## 11. Discussion

### 11.1. What has been achieved?

In this work we have described and developed a fully decision-theoretic approach to the problem of dynamic treatment assignment. The central issue identified and addressed is the transfer of probabilistic information between differing regimes. When justified, this can allow future policy analysis to take appropriate account of previously gathered data.

Out of this approach we have developed an alternative derivation and interpretation of Robins's $G$-computation algorithm, relating it to the fundamental 'backward induction' recursion algorithm of dynamic programming. Moreover we have shown that this is applicable more generally, including to problems involving randomized treatment decisions.

We have devoted some attention to the question of how one might justify the simple stability property (6), or the more general conditions of Lemma 8.1. One can attempt this by including unobservable variables into one's reasoning, and using influence diagram to check the desired properties by simple graphical manipulations. However, as discussed in §7.2.2, the graphical approach sometimes imposes more restrictions than necessary, and an algebraic approach based on manipulations of conditional independence properties can be more general and powerful.

We have also broadened the application of the graphical approach of Pearl and Robins (1995) to allow assessment of the effects of conditional interventions, that are allowed to depend on the values of other variables in the problem. This is a particularly natural requirement when we contemplate sequential interventions, where it is clearly desirable to be able to respond appropriately to the information obtained to date, and so naturally to consider dynamic strategies. We have noted that the graphical expression of condition (6) for simple stability is equivalent to sequential application of Pearl's back-door criterion, and that this allows identification by $G$-recursion of the consequences of conditional interventions, not only for the ultimate response $Y$ but also for every intermediate covariate $L_{i}$. We have further noted that our graphical check for the more general case of $\S 8$ is equivalent to that suggested by Robins (1997).

### 11.2. Syntax and semantics

An important pragmatic aspect of our approach is that, in order to apply it sensibly, we have to be very clear about the real-world meaning of all the vari-
ables (whether 'random' or 'decision') appearing in our formulae. Thus, when considering some interventional regime, we need to understand exactly what real-world interventions are involved: we can not assume that setting a variable to a specific value in different ways, or in different contexts, will have the same overall effects on the system studied - see Hernán and Taubman (2008) for a discussion of these issues in the context of a specific application. Whenever we consider arguments in favour of or against accepting a condition such as stability or extended stability, we must do so in full appreciation of the applied context and circumstances - there can be no purely formal way of addressing such issues.

This emphasis on the semantics of our representations contrasts with that of other popular approaches, such as causal interpretation of DAGs or the docalculus (Pearl, 2009), which appear to operate purely syntactically. However that is an illusion, since those interpretations and manipulations are always grounded in an already assumed formal representation of the problem (e.g. as a DAG, or a set of structural equations). So until we have satisfied ourselves that this representation truly does capture our understanding of the real-world behaviour of our problem - in particular, that it correctly describes the effects of the interventions we care about - there can be no reason to have any faith in the results of any formal manipulations on it.

### 11.3. Statistical inference

We have not directly addressed problems of statistical inference. One might want to estimate the consequences of some proposed sequential strategy, or test a null hypothesis that no strategy is effective in controlling the outcome. In principle one can estimate the ingredients of the $G$-recursion formula, either parametrically or non-parametrically, from the available data, and then (assuming simple stability, or the more general conditions of Lemma 8.1) apply it to supply estimates or tests of the effects of strategies of interest. The proposal by Arjas and Saarela (2010) can be regarded as a Bayesian version of $G$-computation. However, as pointed out by Robins and Wasserman (1997), naïve use of parametric models for the required conditional distributions can lead to a 'null-paradox', rendering it impossible to discover that different strategies have the same consequences. Also, when continuous variables are included, $G$-recursion can involve a large number of nested integrals and become computationally impossible to implement. Hence we find only a few instances where $G$-computation has been used for practical data analysis (Robins, Greenland and $\mathrm{Hu}, 1999$; Taubman et al., 2009). The problems in applying $G$-recursion are exacerbated by the need, in many practical applications, to choose a large set of covariates $\mathcal{L}$ so as to justify the stability assumption. This makes the modelling task more difficult and raises issues of robustness to misspecification. Such considerations have motivated the introduction of marginal or nested 'structural models' (Robins, 1998; Robins, Hernán and Brumback, 2000), as well as doublyrobust methods (Kang and Schafer, 2007), avoiding the null-paradox. Note that
while $G$-recursion provides a likelihood-based approach to the estimation of the consequence of a given strategy, these latter methods rely on estimating equations. It should be straightforward to reinterpret these models and analyses within a fully decision-theoretic framework, by appropriate modelling of the intervention distributions $p(\cdot ; s)$.

### 11.4. Optimal dynamic treatment strategies

Our work is motivated in part by the desire to compare a variety of sequential treatment strategies so as to identify the best one. Recall that our set of regimes is given by $\mathcal{S}=\{o\} \cup \mathcal{S}^{*}$, where $o$ is the observational regime, and $\mathcal{S}^{*}$ is the set of interventional strategies that we want to compare. If we want to apply $G$-recursion, justifying it by simple stability as in $\S 5.3$ or by the more general conditions of Lemma 8.1, we need to ensure that the respective conditions hold for all strategies $e \in \mathcal{S}^{*}$ that we want to compare. As we saw in $\S 8.3 .2$, this is not trivial: if $\mathcal{S}^{*}$ contains static as well as dynamic strategies, in some situations the former may be identified while the latter are not. In fact it follows from Dawid and Didelez (2008) that if want to find an optimal strategy among all dynamic regimes, we will usually need the restrictive requirement of simple stability to hold for all $e \in \mathcal{S}^{*}$.

As mentioned in $\S 4$, the standard dynamic programming routine for identifying an optimal strategy can be regarded as a combination of $G$-recursion and stagewise optimisation. Under conditions allowing $G$-recursion, this can in principle be put directly into effect, after estimating all the required distributional ingredients from the available data. In practice (as pointed out by Robins (1986) and many others since), this quickly becomes infeasible, especially if one wants to avoid parametric restrictions. This is because the number of possible histories for which the optimal next decision has to be determined at each stage of the backward induction recursion can grow extremely rapidly with increasing number $N$ of time points and levels of $\left(\bar{l}_{i}, \bar{a}_{i-1}\right)$.

Alternative approaches to the optimisation problem to sidestep this computational complexity have been suggested. Murphy (2003) introduces a method based on regret functions (see the discussion and application in Rosthøj et al. (2006)), which is closely related to the structural nested models of Robins (2004) (see Moodie, Richardson and Stephens (2007) for a comparison of these two approaches). Henderson, Ansel and Alshibani (2010) modify Murphy's approach so as to be amenable to standard statistical model checking procedures. However, all these alternative methods for finding optimal dynamic treatments rely on the same identification conditions underlying $G$-computation, as well as on various additional (semi-)parametric assumptions.

### 11.5. Complete identifiability

Simple stability, or the alternative conditions of Lemma 8.1, are sufficient conditions allowing the use of $G$-recursion, and thereby identification of the consequences of a given strategy. In recent years the Artificial Intelligence community
has devoted some effort to finding necessary as well as sufficient conditions for the identifiability of consequences of interventions (Huang and Valtorta, 2006; Shpitser and Pearl, 2006a,b). These results rely heavily on the assumptions encoded in causal DAGs or semi-Markovian causal models. Even within this more restricted framework, the general question of identifiability of dynamic treatment strategies seems still to be an open problem (but see Tian (2008)).

### 11.6. Other problems

Many problems in causal inference, previously tackled using potential response or causal DAG formulations, gain in clarity, simplicity and generality when reformulated as problems of decision analysis. Specific topics that have been fruitfully treated in this way include: confounding (Dawid, 2002); partial compliance (Dawid, 2003); direct and indirect effects (Didelez, Dawid and Geneletti, 2006; Geneletti, 2007); identification of the effect of treatment on the treated (Geneletti and Dawid, 2010); Mendelian randomization (Didelez and Sheehan, 2007); Granger causality (Eichler and Didelez, 2010); and causal inference under outcome-dependent sampling (Didelez, Kreiner and Keiding, 2010). However there still remains a wide range of other issues in 'causal inference' that we believe would benefit from the application of the decision-theoretic viewpoint.

## Acknowledgment

We are indebted to Susan Murphy for stimulating this work and for many valuable comments. We also want to thank Jamie Robins for helpful discussions. Financial support from MRC Collaborative Project Grant G0601625 is gratefully acknowledged.

## Appendix A: Two lemmas on DAG-separation

Here we prove generalised versions of equations (8) and (9) (Lemma 1) of Pearl and Robins (1995).

Let $\mathcal{D}$ be a DAG.

## Lemma A. 1

$$
\begin{equation*}
Y \Perp_{\mathcal{D}} X\left|Z \Rightarrow Y \Perp_{\mathcal{D}} X\right| Z^{*} \tag{58}
\end{equation*}
$$

whenever $Z \subseteq Z^{*} \subseteq \operatorname{an}(X \cup Y \cup Z)$.
Proof. Let $\mathcal{G}:=\operatorname{man}(X \cup Y \cup Z)$; then also $\mathcal{G}=\operatorname{man}\left(X \cup Y \cup Z^{*}\right)$. The left-hand side of (58) says that any path from $Y$ to $X$ in $\mathcal{G}$ intersects $Z$, whence it must also intersect the larger set $Z^{*}$.

## Lemma A. 2

$$
\begin{equation*}
Y \Perp_{\mathcal{D}} X\left|Z \Rightarrow Y \Perp_{\mathcal{D}} X\right| Z^{*} \tag{59}
\end{equation*}
$$

whenever $Z^{*}=Z \cap A$ for $A$ an ancestral set in $\mathcal{D}$ containing $X \cup Y$.

Proof. We first note that $(X \cup Y) \cup Z^{*}$ is a subset of $A$, since both its terms are. Since $A$ is ancestral, it follows that

$$
\begin{equation*}
\operatorname{an}\left(X \cup Y \cup Z^{*}\right) \subseteq A \tag{60}
\end{equation*}
$$

Define $\mathcal{G}$ as above, and $\mathcal{G}^{\prime}:=\operatorname{man}\left(X \cup Y \cup Z^{*}\right)$. Then both the node-set and edge-set for $\mathcal{G}^{\prime}$ are subsets of the corresponding sets for $\mathcal{G}$, and hence the same property holds for the path-set. Suppose the right-hand side of (59) fails. Then there exists a path $\pi$ in $\mathcal{G}^{\prime}$ connecting $Y$ and $X$ and avoiding $Z^{*}$; then $\pi$ is a path in $\mathcal{G}$ with the same property. Since $\pi \subseteq \mathcal{G}^{\prime}$, if it intersects $Z$ anywhere it can only do so at a point of $\operatorname{an}\left(X \cup Y \cup Z^{*}\right)$ - and thus, by (60), at a point in $A$, and hence in $Z^{*}$. Since this has been excluded, the result follows.

## References

Arjas, E. and Parner, J. (2004). Causal reasoning from longitudinal data. Scandinavian Journal of Statistics 31 171-187. MR2066247
Arjas, E. and SaArela, O. (2010). Optimal dynamic regimes: Presenting a case for predictive inference. The International Journal of Biostatistics 6. http://tinyurl.com/33dfssf MR2602553
Cowell, R. G., Dawid, A. P., Lauritzen, S. L. and Spiegelhalter, D. J. (1999). Probabilistic Networks and Expert Systems. Springer, New York. MR1697175
DAWID, A. P. (1979). Conditional independence in statistical theory (with Discussion). Journal of the Royal Statistical Society, Series B 41 1-31. MR0535541
Dawid, A. P. (1992). Applications of a general propagation algorithm for probabilistic expert systems. Statistics and Computing 2 25-36.
Dawid, A. P. (1998). Conditional independence. In Encyclopedia of Statistical Science (Update Volume 2) (S. Kotz, C. B. Read and D. L. Banks, eds.) 146-155. Wiley-Interscience, New York.
Dawid, A. P. (2000). Causal inference without counterfactuals (with Discussion). Journal of the American Statistical Association 95 407-448. MR1803167
Dawid, A. P. (2001). Separoids: A mathematical framework for conditional independence and irrelevance. Annals of Mathematics and Artificial Intelligence 32 335-372. MR1859870
Dawid, A. P. (2002). Influence diagrams for causal modelling and inference. International Statistical Review 70 161-189. Corrigenda, ibid., 437.
Dawid, A. P. (2003). Causal inference using influence diagrams: The problem of partial compliance (with Discussion). In Highly Structured Stochastic Systems (P. J. Green, N. L. Hjort and S. Richardson, eds.) 45-81. Oxford University Press. MR2082406
Dawid, A. P. (2010). Beware of the DAG! In Proceedings of the NIPS 2008 Workshop on Causality. Journal of Machine Learning Research Workshop
and Conference Proceedings (D. Janzing, I. Guyon and B. Schölkopf, eds.) 659-86. http://tinyurl.com/33va7tm
Dawid, A. P. and Didelez, V. (2008). Identifying optimal sequential decisions. In Proceedings of the Twenty-Fourth Annual Conference on Uncertainty in Artificial Intelligence (UAI-08) (D. McAllester and A. Nicholson, eds.). 113-120. AUAI Press, Corvallis, Oregon. http://tinyurl.com/3899qpp
Dechter, R. (2003). Constraint Processing. Morgan Kaufmann Publishers.
Didelez, V., Dawid, A. P. and Geneletti, S. G. (2006). Direct and indirect effects of sequential treatments. In Proceedings of the Twenty-Second Annual Conference on Uncertainty in Artificial Intelligence (UAI-06) (R. Dechter and T. Richardson, eds.). 138-146. AUAI Press, Arlington, Virginia. http://tinyurl.com/32w3f4e
Didelez, V., Kreiner, S. and Keiding, N. (2010). Graphical models for inference under outcome dependent sampling. Statistical Science (to appear).
Didelez, V. and Sheehan, N. S. (2007). Mendelian randomisation: Why epidemiology needs a formal language for causality. In Causality and Probability in the Sciences, (F. Russo and J. Williamson, eds.). Texts in Philosophy Series 5 263-292. College Publications, London.
Eichler, M. and Didelez, V. (2010). Granger-causality and the effect of interventions in time series. Lifetime Data Analysis 16 3-32. MR2575937
Ferguson, T. S. (1967). Mathematical Statistics: A Decision Theoretic Approach. Academic Press, New York, London. MR0215390
Geneletti, S. G. (2007). Identifying direct and indirect effects in a noncounterfactual framework. Journal of the Royal Statistical Society: Series B 69 199-215. MR2325272
Geneletti, S. G. and Dawid, A. P. (2010). Defining and identifying the effect of treatment on the treated. In Causality in the Sciences (P. M. Illari, F. Russo and J. Williamson, eds.) Oxford University Press (to appear).

Gill, R. D. and Robins, J. M. (2001). Causal inference for complex longitudinal data: The continuous case. Annals of Statistics 29 1785-1811. MR1891746
Guo, H. and Dawid, A. P. (2010). Sufficient covariates and linear propensity analysis. In Proceedings of the Thirteenth International Workshop on Artificial Intelligence and Statistics, (AISTATS) 2010, Chia Laguna, Sardinia, Italy, May 13-15, 2010. Journal of Machine Learning Research Workshop and Conference Proceedings (Y. W. Teh and D. M. Titterington, eds.) 9 281-288. http://tinyurl.com/331muj7
Henderson, R., Ansel, P. and Alshibani, D. (2010). Regret-regression for optimal dynamic treatment regimes. Biometrics (to appear). doi:10.1111/j.1541-0420.2009.01368.x
Hernán, M. A. and Taubman, S. L. (2008). Does obesity shorten life? The importance of well defined interventions to answer causal questions. International Journal of Obesity 32 S8-S14.
Holland, P. W. (1986). Statistics and causal inference (with Discussion). Journal of the American Statistical Association 81 945-970. MR0867618

Huang, Y. and Valtorta, M. (2006). Identifiability in causal Bayesian networks: A sound and complete algorithm. In AAAI'06: Proceedings of the 21st National Conference on Artificial Intelligence 1149-1154. AAAI Press.
Kang, J. D. Y. and Schafer, J. L. (2007). Demystifying double robustness: A comparison of alternative strategies for estimating a population mean from incomplete data. Statistical Science 22 523-539. MR2420458
Lauritzen, S. L., Dawid, A. P., Larsen, B. N. and Leimer, H. G. (1990). Independence properties of directed Markov fields. Networks 20 491-505. MR1064735
Lok, J., Gill, R., van der Vaart, A. and Robins, J. (2004). Estimating the causal effect of a time-varying treatment on time-to-event using structural nested failure time models. Statistica Neerlandica 58 271-295. MR2157006
Moodie, E. M., Richardson, T. S. and Stephens, D. A. (2007). Demystifying optimal dynamic treatment regimes. Biometrics 63 447-455. MR2370803
Murphy, S. A. (2003). Optimal dynamic treatment regimes (with Discussion). Journal of the Royal Statistical Society, Series B 65 331-366. MR1983752
Oliver, R. M. and Smith, J. Q., eds. (1990). Influence Diagrams, Belief Nets and Decision Analysis. John Wiley and Sons, Chichester, United Kingdom. MR1056324
Pearl, J. (1995). Causal diagrams for empirical research (with Discussion). Biometrika 82 669-710. MR1380809
Pearl, J. (2009). Causality: Models, Reasoning and Inference, Second ed. Cambridge University Press, Cambridge. MR2548166
Pearl, J. and Paz, A. (1987). Graphoids: A graph-based logic for reasoning about relevance relations. In Advances in Artificial Intelligence (D. Hogq and L. Steels, eds.) II 357-363. North-Holland, Amsterdam.
Pearl, J. and Robins, J. (1995). Probabilistic evaluation of sequential plans from causal models with hidden variables. In Proceedings of the 11th Conference on Uncertainty in Artificial Intelligence (P. Besnard and S. Hanks, eds.) 444-453. Morgan Kaufmann Publishers, San Francisco. MR1615028
Raiffa, H. (1968). Decision Analysis. Addison-Wesley, Reading, Massachusetts.
Robins, J. M. (1986). A new approach to causal inference in mortality studies with sustained exposure periods-Application to control of the healthy worker survivor effect. Mathematical Modelling 7 1393-1512. MR0877758
Robins, J. M. (1987). Addendum to "A new approach to causal inference in mortality studies with sustained exposure periods-Application to control of the healthy worker survivor effect". Computers $\xi$ Mathematics with Applications 14 923-945. MR0922792
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 (L. Sechrest, H. Freeman and A. Mulley, eds.) 113-159. NCSHR, U.S. Public Health Service.

Robins, J. M. (1992). Estimation of the time-dependent accelerated failure time model in the presence of confounding factors. Biometrika 79 321-324. MR1185134
Robins, J. M. (1997). Causal inference from complex longitudinal data. In Latent Variable Modeling and Applications to Causality, (M. Berkane, ed.). Lecture Notes in Statistics 120 69-117. Springer-Verlag, New York. MR1601279
Robins, J. M. (1998). Structural nested failure time models. In Survival Analysis, (P. K. Andersen and N. Keiding, eds.). Encyclopedia of Biostatistics 6 4372-4389. John Wiley and Sons, Chichester, UK.
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 6-10.
Robins, J. M. (2004). Optimal structural nested models for optimal sequential decisions. In Proceedings of the Second Seattle Symposium on Biostatistics (D. Y. Lin and P. Heagerty, eds.) 189-326. Springer, New York. MR2129402
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. MR1723276
Robins, J. M., Hernán, M. A. and Brumback, B. (2000). Marginal structural models and causal inference in epidemiology. Epidemiology 11 550-560.
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 (D. Geiger and P. Shenoy, eds.) 409-420. Morgan Kaufmann Publishers, San Francisco. http://tinyurl.com/33ghsas
Rosthøj, S., Fullwood, C., Henderson, R. and Stewart, S. (2006). Estimation of optimal dynamic anticoagulation regimes from observational data: A regret-based approach. Statistics in Medicine 25 4197-4215. MR2307585
Shpitser, I. and Pearl, J. (2006a). Identification of conditional interventional distributions. In Proceedings of the 22nd Annual Conference on Uncertainty in Artificial Intelligence (UAI-06) (R. Dechter and T. Richardson, eds.). 437-444. AUAI Press, Corvallis, Oregon. http://tinyurl.com/2um8w47
Shpitser, I. and Pearl, J. (2006b). Identification of joint interventional distributions in recursive semi-Markovian causal models. In Proceedings of the Twenty-First National Conference on Artificial Intelligence 1219-1226. AAAI Press, Menlo Park, California.
Spirtes, P., Glymour, C. and Scheines, R. (2000). Causation, Prediction and Search, Second ed. Springer-Verlag, New York.
Sterne, J. A. C., May, M., Costagliola, D., de Wolf, F., Phillips, A. N., Harris, R., Funk, M. J., Geskus, R. B., Gill, J., Dabis, F., Miro, J. M., Justice, A. C., Ledergerber, B., Fatkenheuer, G., Hogg, R. S., D'Arminio-Monforte, A., Saag, M., Smith, C., Staszewski, S., Egger, M., Cole, S. R. and When To

Start Consortium (2009). Timing of initiation of antiretroviral therapy in AIDS-Free HIV-1-infected patients: A collaborative analysis of 18 HIV cohort studies. Lancet 373 1352-1363.
taubman, S. L., Robins, J. M., Mittleman, M. A. and Hernán, M. A. (2009). Intervening on risk factors for coronary heart disease: An application of the parametric $g$-formula. International Journal of Epidemiology 38 15991611.

Tian, J. (2008). Identifying dynamic sequential plans. In Proceedings of the Twenty-Fourth Annual Conference on Uncertainty in Artificial Intelligence (UAI-08) (D. McAllester and A. Nicholson, eds.). 554-561. AUAI Press, Corvallis, Oregon. http://tinyurl.com/36ufx2h
Verma, T. and Pearl, J. (1990). Causal networks: Semantics and expressiveness. In Uncertainty in Artificial Intelligence 4 (R. D. Shachter, T. S. Levitt, L. N. Kanal and J. F. Lemmer, eds.) 69-76. North-Holland, Amsterdam. MR1166827


[^0]:    *This paper was accepted by Elja Arjas, Executive Editor for the Bernoulli.

[^1]:    ${ }^{1}$ Realistically the loss could also depend on the values of intermediate variables, e.g. if these relate to adverse drug reactions. Such problems can be treated by redefining $Y$ as the overall loss suffered (at any rate so long as this loss does not depend on other, unobserved, variables.)
    ${ }^{2}$ More correctly, these correspond to what are termed behavioral rules in decision theory (Ferguson, 1967)

[^2]:    ${ }^{3}$ More generally (see footnote 1 ), we could consider a function $Y^{*}$ of ( $\bar{L}_{N}, \bar{A}_{N}, Y$ ). Starting now with $f\left(\bar{l}_{N}, \bar{a}_{N}, y\right):=Y^{*}\left(\bar{l}_{N}, \bar{a}_{N}, y\right)$, we can apply the identical steps to arrive at $f(\emptyset)=$ $\mathrm{E}\left\{Y^{*} ; s\right\}$. In particular we can evaluate the expected overall loss under $s$, even when the loss function depends on the full sequence of variables.

[^3]:    ${ }^{4}$ Cases in which simple stability may not hold but we can nevertheless still apply $G$ recursion are considered in Section 8.

[^4]:    ${ }^{5}$ An alternative, and entirely equivalent, approach can be based on the ' $d$-separation criterion' (Pearl, 2009; Verma and Pearl, 1990). We have found (31) more straightforward to understand and apply.

[^5]:    ${ }^{6}$ Both these papers refer for the details to an unpublished paper, Robins and Pearl (1996).

[^6]:    ${ }^{7}$ Note, as a matter of logic, that if we follow $e$ we shall not be able to observe e.g. $Y_{o}$ (though see the note after Condition 10.2 below). This is a version of the so-called 'fundamental problem of causal inference' (Holland, 1986) which has been critically discussed by Dawid (2000).

[^7]:    ${ }^{8}$ This is sometimes expressed in a stronger form that drops the restriction to future variables, so replacing $L_{e}^{i+1}$ by $\left(\bar{L}_{e, N}, Y_{e}\right)$ (Robins, 2000).

