Chapter 2
The Data: Observational Studies and Sequentially Randomized Trials

The data for constructing (optimal) DTRs that we consider are obtained from either longitudinal observational studies or sequentially randomized trials. In this chapter we review these two types of data sources, their advantages and drawbacks, and the assumptions required to perform valid analyses in each, along with some examples. We also discuss a basic framework of causal inference in the context of observational studies, and power and sample size issues in the context of randomized studies.

2.1 Longitudinal Observational Studies

The goal of much of statistical inference is to quantify causal relationships, for instance to be able to assert that a specified treatment[1] improves patient outcomes rather than to state that treatment use or prescription of treatment is merely associated or correlated with better patient outcomes. Randomized trials are the "gold standard" in study design, as randomization coupled with compliance allows causal interpretations to be drawn from statistical association. Making causal inferences from observational data, however, can be tricky and relies critically on certain (unverifiable) assumptions which we will discuss in Sect. 2.1.3. The notion of causation is not new: it has been the subject matter of philosophers as far back as Aristotle, and more recently of econometricians and statisticians. Holland (1986) provides a nice overview of the philosophical views and definitions of causation as well as of the causal models frequently used in statistics. Neyman (1923) and later Rubin (1974) laid the foundations for the framework now used in modern causal inference. The textbook *Causal Inference* (Hernán and Robins 2013) provides a thorough description of basic definitions and most modern methods of causal inference for both

[1] In this book, we use the term *treatment* generically to denote either a medical treatment or an *exposure* (which is the preferred term in the causal inference literature and more generally in epidemiology).

B. Chakraborty and E.E.M. Moodie, *Statistical Methods for Dynamic Treatment Regimes*, Statistics for Biology and Health 76, DOI 10.1007/978-1-4614-7428-9_2,

point-source treatment (i.e. cross-sectional, or one-stage) settings as well as general longitudinal settings with time-varying treatments and the associated complexities.

2.1.1 The Potential Outcomes Framework

Much of the exposition of methods used when data are observational will rely on the notion of *potential outcomes* (also called *counterfactuals*), defined as a person's outcome had he followed a particular treatment regime, possibly different from the regime which he was actually observed to follow (hence, counter to fact). The individual-level causal effect of a regime may then be viewed as the difference in outcomes if a person had followed that regime as compared to a placebo regime or a standard care protocol. Consider, for example, a simple one-stage[2] randomized trial in which subjects can receive either a or a'. Suppose now that an individual was randomized to receive treatment a. This individual will have a single observed outcome Y which corresponds to the potential outcome "Y under treatment a", denoted by $Y(a)$, and one unobservable potential outcome, $Y(a')$, corresponding to the outcome under a'. An alternative notation to express counterfactual quantities is via subscripting: Y_a and $Y_{a'}$ (Hernán et al. 2000). Pearl (2009) uses an approach similar to that of the counterfactual framework, using what is called the "do" notation to express the idea that a treatment is administered rather than simply observed to have been given: in his notation, $E[Y|do(A = a)]$ is the expected value of the outcome variable Y under the intervention regime a, i.e. it is the population average were all subjects forced to take treatment a.

The so-called *fundamental problem of causal inference* lies in the definition of causal parameters at an individual level. Suppose we are interested in the causal effect of taking treatment a instead of treatment a'. An individual-level causal parameter that could be considered is a person's outcome under treatment a' subtracted from his outcome under treatment a, i.e. $Y(a) - Y(a')$. Clearly, it is not possible to observe the outcome under both treatments a and a' without further data and assumptions (e.g. in a cross-over trial with no carry-over effect) and so the individual-level causal effect can never be observed. However, population-level causal parameters or average causal effects can be identified under randomization with perfect compliance, or bounded under randomization with non-compliance. Without randomization, i.e. in observational studies or indeed randomized trials with imperfect compliance, further assumptions are required to estimate population-level causal effects, which we shall detail shortly.

Suppose now that rather than being a one-stage trial, subjects are treated over two stages, and can receive at each stage either a or a'. If an individual was randomized to receive treatment a first and then treatment a', this individual will have a single observed outcome Y which corresponds to the potential outcome "Y under regime

[2] While the term *stage* is commonly used in the randomized trial literature, the term *interval* is more popular in the causal inference literature. In this book, for consistency, we will use the term *stage* for both observational and randomized studies.

(a,a')", which we denote by $Y(a,a')$, and three unobservable potential outcomes: $Y(a,a)$, $Y(a',a)$, and $Y(a',a')$, corresponding to outcomes under each of the other three possible regimes. As is clear even in this very simple example, the number of potential outcomes and causal effects as represented by contrasts between the potential outcomes can be very large, even for a moderate number of stages. As shall be seen in Chap. 4, the optimal dynamic regime may be estimated while limiting the models specified to only a subset of all possible contrasts.

2.1.2 Time-Varying Confounding and Mediation

Longitudinal data are increasingly available to health researchers; this type of data presents challenges not observed in cross-sectional data, not the least of which is the presence of time-varying confounding variables and intermediate effects. A variable O is said to be a *mediating* or *intermediate* variable if it is caused by A and in turn causes changes in Y. For example, a prescription sleep-aid medication (A) may cause dizziness (O) which in turn causes fall-related injuries (Y). In contrast, a variable, O, is said to *confound* a relationship between a treatment A and an outcome Y if it is a common cause of both the treatment and the outcome. More generally, a variable is said to be a confounder (relative to a set of covariates X) if it is a pre-treatment covariate that removes some or all of the bias in a parameter estimate, when taken into account in addition to the variables X. It may be the case, then, that a variable is a confounder relative to one set of covariates X but not another, X'. If the effect of O on both A and Y is not accounted for, it may appear that there is a relationship between A and Y when in fact their pattern of association may be due entirely to changes in O. For example, consider a study of the dependence of the number of deaths by drowning (Y) on the use of sunscreen (A). A strong positive relationship is likely to be observed, however it is far more likely that this is due to the confounding variable air temperature (O). When air temperature is high, individuals may be more likely to require sunscreen and may also be more likely to swim, but there is no reason to believe that the use of sunscreen increases the risk of drowning. In cross-sectional data, eliminating the bias due to a confounding effect is typically achieved by adjusting for the variable in a regression model.

Directed Acyclic Graphs (DAGs), also called *causal graphs*, formalize the causal assumptions that a researcher may make regarding the variables he wishes to analyze. A graph is said to be *directed* if all inter-variable relationships are connected by arrows indicating that one variable causes changes in another and *acyclic* if it has no closed loops (no feedback between variables); see, for example, Greenland et al. (1999) or Pearl (2009) for further details. DAGs are becoming more common in epidemiology and related fields as researchers seek to clarify their assumptions about hypothesized relationships and thereby justify modeling choices (e.g. Bodnar et al. 2004; Brotman et al. 2008). In particular, confounding in its simplest form can be visualized in a DAG if there is an arrow from O into A, and another from O into Y. Similarly, mediation is said to occur if there is at least one directed path of arrows from A to Y that passes through O.

Let us now briefly turn to a two-stage setting where data are collected at three time-points: baseline (t_1=0), t_2, and t_3. Covariates are denoted O_1 and O_2, measured at baseline and t_2, respectively. Treatment at stages 1 and 2, received in the intervals $[0,t_2)$ and $[t_2,t_3)$, are denoted A_1 and A_2 respectively. Outcome, measured at t_3, is denoted Y. Suppose there is an additional variable, U, which is a cause of both O_2 and Y. See Fig. 2.1.

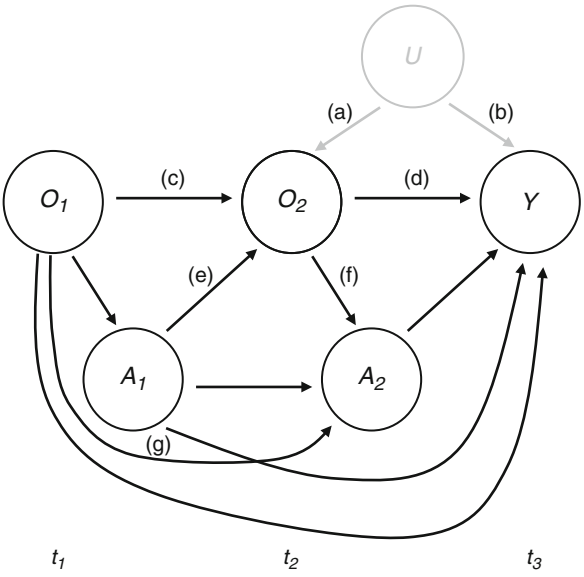

Fig. 2.1 A two-stage directed acyclic graph illustrating time-varying confounding and mediation

We first focus on the effect of A_1 on Y; A_1 acts directly on Y, but also acts indirectly through O_2 as indicated by arrows (e) and (d); O_2 is therefore a mediator. We now turn our attention to the effect of A_2 on Y; O_2 confounds this relationship, as can be observed by arrows (d) and (f). In this situation, adjustment for O_2 is essential to obtaining unbiased estimation of the effect of A_2 on Y. However, complications may arise if there are unmeasured factors that also act as confounders; in Fig. 2.1, U acts in this way. If one were to adjust for O_2 in a regression model, it would open what is called a "back-door" path from Y to A_2 via the path (b)→(a)→(c)→(g). This is known as *collider-stratification bias, selection bias, Berksonian bias, Berkson's paradox*, or, in some contexts, the *null paradox* (Robins and Wasserman 1997; Gail and Benichou 2000; Greenland 2003; Murphy 2005a); this problem will be considered in greater depth in Sect. 3.4.2 in the context of estimation. Collider-stratification bias can also occur when conditioning on or stratifying by variables that are caused by both the exposure and the outcome, and there has been a move in the epidemiology literature to use the term selection bias only for bias caused by conditioning on post-treatment variables, and the term confounding for bias caused by pre-treatment variables (Hernán et al. 2004).

Modeling choices become more complex when data are collected over time, particularly as a variable may act as both a confounder and a mediator. The use of a DAG forces the analyst to be explicit in his modeling assumptions, particularly as the absence of an arrow between two variables ("nodes") in a graph implies the assumption of (conditional) independence. Some forms of estimation are able to avoid the introduction of collider-stratification bias by eliminating conditioning (e.g. weighting techniques) while others rely on the assumption that no variables such as U exist. See Sect. 3.4.2 for a discussion on how *Q-learning*, a stage-wise regression based method of estimation, avoids this kind of bias by analyzing one stage at a time.

2.1.3 Necessary Assumptions

A fundamental requirement of the potential outcomes framework is the *axiom of consistency*, which states that the potential outcome under the observed treatment and the observed outcome agree: that is, the treatment must be defined in such a way that it must be possible for all treatment options to be assigned to all individuals in the population under consideration. Thus, the axiom of consistency requires that the outcome for a given treatment is the same, regardless of the manner in which treatments are 'assigned'. This is often plausible in studies of medical treatments where it is easy to conceive of how to manipulate the treatments given to the patients (this setting is relevant in the DTR context), but less obvious for exposures that are modifiable by a variety of means, such as body-mass index (Hernán and Taubman 2008), or that are better defined as (non-modifiable) characteristics, such as sex (Cole and Frangakis 2009).

Before stating the necessary assumptions for estimating DTRs, we introduce the following notations. Let $\bar{a}_K \equiv (a_1, \ldots, a_K)$ denote a K-stage sequence of treatments. Let $(d_1, \ldots, d_K)$ denote a treatment regime, i.e. a set of decision rules where d_j is a mapping from the history space to the treatment/action space for all j. Similarly let $\bar{O}_j \equiv (O_1, \ldots, O_j)$ denote the collection of covariates observed up to stage j and $\bar{A}_{j-1} \equiv (A_1, \ldots, A_{j-1})$ denote the collection of past treatments prior to stage j. We combine the treatment and covariate history up to the jth stage into a single *history* vector, $H_j \equiv (\bar{O}_j, \bar{A}_{j-1})$. To estimate a DTR from either randomized or observational data, two assumptions are required:

1. *Stable unit treatment value assumption (SUTVA)*: A subject's outcome is not influenced by other subjects' treatment allocation (Rubin 1980).
2. *No unmeasured confounders* (NUC): For any regime $\bar{a}_K$,

$$A_j \perp (O_{j+1}(\bar{a}_j), \ldots, O_K(\bar{a}_{K-1}), Y(\bar{a}_K)) \Big| H_j \quad \forall j = 1, \ldots, K.$$

That is, for any possible regime $\bar{a}_K$, treatment A_j received in the jth stage is independent of any future (potential) covariate or outcome, $O_{j+1}(\bar{a}_j), \ldots, O_K(\bar{a}_{K-1})$, $Y(\bar{a}_K)$, conditional on the history H_j (Robins 1997).

The first assumption – sometimes called *no interaction between units* or *no interference between units* (Cox 1958) – is often reasonable, particularly in the context of randomized trials where study participants are drawn from a large population. SUTVA may be violated in special cases such as vaccinations for contagious disease where the phenomenon of "herd immunity" may lead to protection of unvaccinated individuals or in the context of group therapy (e.g. a support group) where the inter-personal dynamics between group members could influence outcomes.

The NUC assumption always holds under either complete or sequential randomization, and is sometimes called the *sequential randomization assumption* (SRA), *sequential ignorability*, or *exchangeability*, which is closely linked to the concept of *stability* (Dawid and Didelez 2010; Berzuini et al. 2012). The assumption may also be (approximately) true in observational settings where all relevant confounders have been measured. No unmeasured confounding is a strong generalization of the usual concept of randomization in a single-stage trial, whereby it is assumed that, conditional on treatment and covariate history, *at each stage* the treatment actually received, A_j, is independent of future states and outcome under *any* sequence of future treatments, $\bar{a}_j$. That is, conditional on the past history, treatment received at stage j is independent of future potential covariates and outcome:

$$p(A_j|H_j, O_{j+1}(\bar{a}_j), \ldots, O_K(\bar{a}_{K-1}), Y(\bar{a}_K)) = p(A_j|H_j).$$

It is this assumption that allows us to effectively view each stage as a randomized trial, possibly with different randomization probabilities at stage j, given strata defined by the history H_j.

If subjects are censored (lost to follow-up or otherwise removed from the study), we must further assume that censoring is non-informative conditional on history, i.e. that the potential outcomes of those subjects who are censored follow the same distribution as that of those who are fully followed given measured covariates.

The optimal regime may only be estimated non-parametrically among the set of *feasible* regimes (Robins 1994). Let $p_j(a_j|H_j)$ denote the conditional probability of receiving treatment a_j given H_j, and let $f(H_K)$ denote the density function of H_K. Then for all histories h_K with $f(h_K) > 0$, a feasible regime $\bar{d}_K$ satisfies

$$\prod_{j=1}^{K} p_j(d_j(H_j)|H_j = h_j) > 0.$$

That is, feasibility requires some subjects to have followed regime $\bar{d}_K$ for the analyst to be able to estimate its performance non-parametrically. To express this in terms of decision trees, no non-parametric inference can be made about the effect of following a particular branch of a decision tree if no one in the sample followed that path.

Other terms have been used to describe feasible treatment regimes, including *viable* (Wang et al. 2012) and *realistic* (Petersen et al. 2012) rules. Feasibility is closely related to the *positivity*, or *experimental treatment assignment* (ETA), assumption. Positivity, like feasibility, requires that there are both treated and

untreated individuals at every level of the treatment and covariate history. Positivity may be violated either *theoretically* or *practically*. A theoretical or structural violation occurs if the study design prohibits certain individuals from receiving a particular treatment, e.g. failure of one type of drug may preclude the prescription of other drugs in that class. A practical violation of the positivity assumption is said to occur when a particular stratum of subjects has a very low probability of receiving the treatment (Neugebauer and Van der Laan 2005; Cole and Hernán 2008). Visual and bootstrap-based approaches to diagnosing positivity violations have been proposed for one-stage settings (Wang et al. 2006; Petersen et al. 2012). Practical positivity violations may be more prevalent in longitudinal studies if there exists a large number of possible treatment paths; methods for handling such violations in multi-stage settings are less developed.

There is an additional assumption that is not required for estimation, but that is useful for understanding the counterfactual quantities and models that will be considered: the assumption of *additive local rank preservation*, which we shall elucidate in two steps. First, *local rank preservation* states that the ranking of subjects' outcomes under a particular treatment pattern $\overline{a}_K$ is the same as their ranking under any other pattern, say $\overline{d}_K$, given treatment and covariate history (see Table 2.1). In particular, if we consider two regimes $\overline{d}_K$ and $\overline{a}_K$, local rank preservation states that the ranking of patients' outcomes under regime $\overline{d}_K$ is the same as their ranking under regime $\overline{a}_K$ conditional on the history H_j. Local rank preservation is said to be *additive* when $Y(\overline{d}_K) = Y(\overline{a}_K) + \text{cons}$, where $\text{cons} = E[Y(\overline{d}_K) - Y(\overline{a}_K)]$, i.e., the individual causal effect equals the *average causal effect*. This is also called *unit treatment additivity*. Thus, rank preservation makes the assumption that the individuals who do best under one regime will also do so under another, and in fact the ranking of each individual's outcome will remain unchanged whatever the treatment pattern received. Additive local rank preservation makes the much stronger assumption that the difference between any two individuals' outcomes will be the same under all treatment patterns.

Table 2.1 Local rank preservation (LRP) and additive LRP, assuming all subjects have the same baseline covariates

			LRP		Additive LRP	
Subject	$Y(\overline{a}_K)$	Rank	$Y(\overline{d}_K)$	Rank	$Y(\overline{d}_K)$	Rank
1	12.8	3	15.8	3	13.9	3
2	10.9	1	14.0	1	13.0	1
3	13.1	4	16.0	4	14.2	4
4	12.7	2	14.5	2	13.8	2

2.2 Examples of Longitudinal Observational Studies

A variety of studies aimed at estimating optimal DTRs from observational data have been undertaken. Data sources include administrative (e.g. hospital) databases (Rosthøj et al. 2006; Cain et al. 2010; Cotton and Heagerty 2011), randomized

encouragement trials (Moodie et al. 2009), and cohort studies (Van der Laan and Petersen 2007b). We shall briefly describe three here to demonstrate the variety of questions that can be addressed using observational data and DTR methodology. In particular, the data in the examples below have been addressed using regret-regression, G-estimation, and marginal structural models; these and related methods of estimation are presented in Chaps. 4 and 5.

2.2.1 Investigating Warfarin Dosing Using Hospital Data

Rosthøj et al. (2006) aimed to find a warfarin dosing strategy to control the risk of both clotting and excessive bleeding, by tailoring treatment using the international normalized ratio, a measure of clotting tendency of blood. Observational data were taken from hospital records over a five year period; recorded variables included age, sex, and diagnosis as well as a time-varying measure of INR. There exists a standard target range for INR, and so the vector-valued tailoring variable, O_j, was taken to be 0 if the most recent INR measurement lay within the target range and otherwise was taken to be the ratio of the difference between the INR measurement and the nearest boundary of the target range, and the width of that target range. Treatment at stage j, A_j, was taken to be the change in warfarin dose (with 0 being an acceptable option). The outcome of interest was taken to be the percentage of the time on study in which a subject's INR was within the target range.

Rosthøj et al. (2006) modeled the effect of taking the observed rather than the optimal dose of warfarin using parametric mean models that are quadratic in the dosing effect so that doses that are either too low or too high are penalized.

2.2.2 Investigating Epoetin Therapy Using the United States Renal Data System

Cotton and Heagerty (2011) performed an analysis of the United States Renal Data System, an administrative data set based on Medicare claims for hemodialysis with end-stage renal disease. Covariates included demographic variables as well as clinical and laboratory variables such as diabetes, HIV status, and creatinine clearance. Monthly information was also available on the number of dialysis sessions reported, the number of epoetin doses recorded, the total epoetin dosage, iron supplementation dose, the number of days hospitalized and the most recently recorded hematocrit measurement in the month.

Restricting their analysis to incident end-stage renal disease patients free from HIV/AIDS from 2003, Cotton and Heagerty (2011) considered treatment rules that adjust epoetin treatment at time j, A_j, multiplicatively based on the value of treatment in the previous month, A_{j-1}, and the most recent hematocrit measurement, O_j:

$$A_j \in \begin{cases} A_{j-1} \times (0, 0.75) & \text{if } O_j \geq \psi - 3 \\ A_{j-1} \times (0.75, 1.25) & \text{if } O_j \in (\psi - 3, \psi + 3) \\ A_{j-1} \times (1.25, \infty) & \text{if } O_j \leq \psi + 3 \end{cases}$$

where the target hematocrit range specified by the parameter ψ is varied to consider a range of different regimes. That is, O_j is the tailoring variable at each month, and the optimal regime is the treatment rule $d_j^{opt}(O_j, A_{j-1}; \psi)$ that maximizes survival time for $\psi \in \{31, 32, \ldots, 40\}$. Thus, in contrast to the strategy employed by Rosthøj et al. (2006), the decision rules considered in the analysis of Cotton and Heagerty (2011) did not attempt to estimate the optimal treatment changes/doses, but rather focused on estimating which target range of hematocrit should initiate a change in treatment dose from one month to the next. Note that the parameter ψ (the mid-value of the target hematocrit range) does not vary over time, but rather is common over all months; this is called *parameter sharing* (over time).

2.2.3 Estimating Optimal Breastfeeding Strategies Using Data from a Randomized Encouragement Trial

The Promotion of Breastfeeding Intervention Trial (PROBIT) (Kramer et al. 2001) has been used to explore several different dynamic regimes, with a view to optimizing growth (Moodie et al. 2009; Rich et al. 2010) and the vocabulary subtest of the Wechsler Abbreviated Scales of Intelligence (Moodie et al. 2012).

PROBIT randomized hospitals and affiliated polyclinics in the Republic of Belarus to a breastfeeding promotion intervention modeled on the WHO/UNICEF Baby-Friendly Hospital Initiative or standard care. Mother-infant pairs were enrolled during their postpartum stay, and follow-up visits were scheduled at 1, 2, 3, 6, 9, and 12 months of age for various measures of health and size, including weight, length, number of hospitalizations and gastrointestinal infections since the last scheduled visit. At each follow-up visit up to 12 months, it was established whether the infant was breastfeeding, as well as whether the infant was given other liquids or solid foods. In a later wave of PROBIT, follow-up interviews and examinations including the Wechsler test were performed on 13,889 (81.5 %) children at 6.5 years of age.

In analyses of these data, the treatment A_j was taken to be continued breastfeeding throughout the jth stage, and variables such as infant weight at the start of the stage or the number of gastrointestinal infections at the previous stage have been considered as potential tailoring variables, O_j.

2.3 Sequentially Randomized Studies

It is well known that estimates based on observational data are often subject to *confounding* and various hidden biases; hence randomized data, when available, are preferable for more accurate estimation and stronger statistical inference (Rubin 1974; Holland 1986; Rosenbaum 1991). This is especially important when dealing with DTRs since the hidden biases can compound over stages. One crucial point to note here is that developing DTRs is a developmental procedure rather than a confirmatory procedure. Usual randomized controlled trials are used as the "gold standard" for evaluating or confirming the efficacy of a newly developed treatment, not for developing the treatment *per se*. Thus, generating meaningful data for developing optimal DTRs is beyond the scope of the usual confirmatory randomized trials; special design considerations are required. A special class of designs called *sequential multiple assignment randomized trial* (SMART) designs, tailor-made for the purpose of developing optimal DTRs, is discussed below.

SMART designs involve an initial randomization of patients to possible treatment options, followed by re-randomizations at each subsequent stage of some or all of the patients to another treatment available at that stage. The re-randomizations at each subsequent stage may depend on information collected after previous treatments, but prior to assigning the new treatment, e.g. how well the patient responded to the previous treatment. Thus, even though a subject is randomized more than once, ethical constraints are not violated. This type of design was first introduced by Lavori and Dawson (2000) under the name *biased coin adaptive within-subject* (BCAWS) design, and practical considerations for designing such trials were discussed by Lavori and Dawson (2004). Building on these works, Murphy (2005a) proposed the general framework of the SMART design. These designs attempt to conform better to the way clinical practice for chronic disorders actually occurs, but still retain the well-known advantages of randomization over observational studies.

SMART-like trials, i.e. trials involving multiple randomizations had been used in various fields even before the exact framework was formally established; see for example, the CALGB Protocol 8923 for treating elderly patients with leukemia (Stone et al. 1995; Wahed and Tsiatis 2004, 2006), the CATIE trial for antipsychotic medications in patients with Alzheimer's disease (Schneider et al. 2001), the STAR*D trial for treatment of depression (Lavori et al. 2001; Rush et al. 2004; Fava et al. 2003), and some cancer trials conducted at the MD Anderson Cancer Center (Thall et al. 2000). Other examples include a smoking cessation study conducted by the Center for Health Communications Research at the University of Michigan (Strecher et al. 2008; Chakraborty et al. 2010), and a trial of neurobehavioral treatments for patients with metastatic malignant melanoma (Auyeung et al. 2009). More recently, Lei et al. (2012) discussed four additional examples of SMARTs: the Adaptive Characterizing Cognition in Nonverbal Individuals with Autism (CCNIA) Developmental and Augmented Intervention (Kasari 2009) for school-age, nonverbal children with autism spectrum disorders; the Adaptive Pharmacological and Behavioral Treatments for children with attention deficit hyperactivity disorder (ADHD) (see for example, Nahum-Shani et al. 2012a,b); the Adaptive Reinforcement-Based

Treatment for Pregnant Drug Abusers (RBT) (Jones 2010); and the ExTENd study for alcohol-dependent individuals (Oslin 2005). Lei et al. (2012) also discussed the subtle distinctions between different types of SMARTs in terms of the extent of multiple randomizations: (i) SMARTs in which only the non-responders to one of the initial treatments are re-randomized (e.g. CCNIA); (ii) SMARTs in which non-responders to all the initial treatments are re-randomized (e.g. the ADHD trial); and (iii) SMARTs in which both responders and non-responders to all the initial treatments are re-randomized (e.g. RBT, ExTENd).

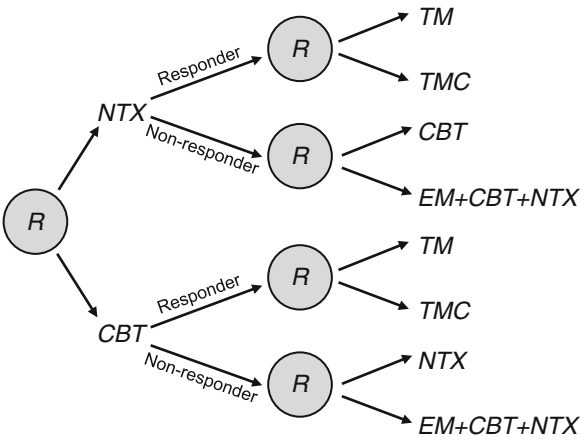

Fig. 2.2 Hypothetical SMART design schematic for the addiction management example (an "R" within a circle denotes randomization at a critical decision point)

In order to make the discussion more concrete, let us consider a hypothetical SMART design based on the addiction management example introduced in Chap. 1; see Fig. 2.2 for a schematic. In this trial, each subject is randomly assigned to one of two possible initial treatments: cognitive behavioral therapy (CBT) or naltrexone (NTX). A subject is classified as a *non-responder* or *responder* to the initial treatment according to whether he does or does not experience more than two heavy drinking days during the next two months. A non-responder to NTX is re-randomized to one of the two subsequent treatment options: either a switch to CBT, or an augmentation of NTX with CBT and an enhanced motivational program (EM + CBT + NTX). Similarly, a non-responder to CBT is re-randomized to either a switch to NTX or an augmentation (EM + CBT + NTX). Responders to the initial treatment are re-randomized to receive either telephone monitoring only (TM) or telephone monitoring and counseling (TMC) for an additional period of six months. The goal of the study is to maximize the number of non-heavy drinking days over a 12-month study period.

2.3.1 SMART Versus a Series of Single-stage Randomized Trials

Note that the goal of SMART design is to generate high quality data that would aid in the development and evaluation of optimal DTRs. A competing design approach could be to conduct separate randomized trials for each of the separate stages, to find the optimal treatment at each stage based on the trial data, and then combine these optimal treatments from individual stages to create a DTR. For example, instead of the SMART design for the addiction management study described above, the researcher may conduct two single-stage randomized trials. The first trial would involve a comparison of the initial treatments (CBT versus NTX). The researcher would then choose the best treatment based on the results of the first trial and move on to the second trial where all subjects would be initially treated with the chosen treatment and then responders would be randomized to one of the two possible options: TM or TMC, and non-responders would be randomized to one of the two possible options: switch of the initial treatment or a treatment augmentation (EM + CBT + NTX). However, when used to optimize DTRs, this approach suffers from several disadvantages as compared to a SMART design.

First, this design strategy is myopic, and may often fail to detect possible *delayed effects* of treatments, ultimately resulting in a suboptimal DTR (Lavori and Dawson 2000). Many treatments can have effects that do not occur until after the intermediate outcome (e.g. response to initial treatment) has been measured, such as improving the effect of a future treatment or long-term side effects that prevent a patient from being able to use an alternative useful treatment in future. SMART designs are capable of taking care of this issue while the competing approach is not. This point can be further elucidated using the addiction management example, following the original arguments of Murphy (2005a). Suppose counseling (TMC) is more effective than monitoring (TM) among responders to CBT; this is a realistic scenario since the subject can learn to use counseling during CBT at the initial stage and thus is able to take advantage of the counseling offered at the subsequent stage to responders. Individuals who received NTX during the initial treatment would not have learned to use counseling, and thus among responders to NTX the addition of counseling to the monitoring does not improve abstinence relative to monitoring alone. If an individual is a responder to CBT, it is best to offer TMC as the secondary treatment. But if the individual is a responder to NTX, it is best to offer the less expensive TM as the secondary treatment. In summary, even if CBT and NTX result in the same proportion of responders (or, even if CBT appears less effective at the initial stage), CBT may be the best initial treatment as part of the entire treatment sequence. This would be due to the enhanced effect of TMC when preceded by CBT. On the other hand, if the researcher employs two separate stage-specific trials, he would likely conduct the second trial with NTX (which is cheaper than CBT) as the initial treatment, unless CBT looks significantly better than NTX at the first trial. In that case, there would be no way for the researcher to discover the truly optimal regime.

Second, even though the results of the first trial may indicate that treatment a is initially less effective than treatment a', it is quite possible that treatment a may elicit

valuable *diagnostic* information that would permit the researcher to better personalize the subsequent treatment to each subject, and thus improve the primary outcome. This issue can be better discussed using the ADHD study example (Nahum-Shani et al. 2012a,b), following the original discussion of Lei et al. (2012). In secondary analyses of the ADHD study, Nahum-Shani et al. (2012a,b) found evidence that children's adherence to the initial intervention could be used to better match the secondary intervention. More precisely, among non-responders to the initial intervention (either low-dose medication or low-dose behavioral modification), those with low adherence performed better when the initial intervention was augmented with the other type of intervention at the second stage, compared to increasing the dose or intensity of the initial treatment at the second stage. This phenomenon is sometimes called the *diagnostic effect* or *prescriptive effect*.

Third, subjects who enroll and remain in a single-stage trial may be inherently different from those who enroll and remain in a SMART. This is a type of *cohort effect* or *selection effect*, as discussed by Murphy et al. (2007a). Consider a single-stage randomized trial in which CBT is compared with NTX. First, in order to reduce variability in the treatment effect, investigators would tend to set very restrictive entry criteria (this is the case with most RCTs), which would result in a cohort that represents only a small subset of the treatable population. In contrast, researchers employing a SMART design would not try to reduce the variability in the treatment effect, since this design would allow varying treatment sequences for different types of patients. Hence SMARTs can recruit from a wider population of patients, and would likely result in greater *generalizability*. Furthermore, in a single-stage RCT, for subjects with no improvement in symptoms and for those experiencing severe side-effects, there is often no option but to drop out of the study or cease to comply with the study protocol. In contrast, non-responding subjects in a SMART would know that their treatments will be altered at some point. Thus it can be argued that non-responding subjects may be less likely to drop out from a SMART relative to a single-stage randomized trial. Consequently the choice of the best initial treatment obtained from a single-stage trial may be based on a sample less representative of the study population compared to the choice of the best initial treatment obtained from a SMART.

From the above discussion, it is clear that conducting separate stage-specific trials and combining best treatment options from these separate trials may fail to detect delayed effects and diagnostic effects, and may result in possible cohort effects, thereby rendering the developed sequence of treatment decisions potentially suboptimal. This has been the motivation to consider SMART designs.

2.3.2 Design Properties

For simplicity of exposition, let us focus on SMART designs with only two stages; however the ideas can be generalized to any finite number of stages. Denote the observable data trajectory for a subject in a SMART by (O_1, A_1, O_2, A_2, Y), where

O_1 and O_2 are the pretreatment information and intermediate outcomes, A_1 and A_2 are the randomly assigned initial and secondary treatments, and Y is the primary outcome, respectively. For example, in the addiction management study discussed earlier, O_1 may include addiction severity and co-morbid conditions, O_2 may include the subject's binary response status, side effects and adherence to the initial treatment, and Y may be the number of non-heavy drinking days over the 12-month study period. Under *the axiom of consistency* (see Sect. 2.1.3), the potential outcomes are connected to the observable data by $O_2 = O_2(A_1)$ and $Y = Y(A_1, A_2)$.

In a SMART, the randomization probabilities may depend on the available treatment and covariate *history*; more precisely, the randomization probabilities for A_1 and A_2 may depend on $H_1 \equiv O_1$ and $H_2 \equiv (O_1, A_1, O_2)$, respectively. Thus data from a SMART satisfy the *sequential ignorability* or *no unmeasured confounding* assumption (see Sect. 2.1.3). Under this assumption, the conditional distributions of the potential outcomes are the same as the corresponding conditional distributions of the observable data. That is,

$$P(O_2(a_1) \leq o_2 | O_1 = o_1) = P(O_2 \leq o_2 | O_1 = o_1, A_1 = a_1),$$

and

$$\begin{aligned} P(Y(a_1, a_2) \leq y | O_1 = o_1, O_2(a_1) = o_2) \\ = P(Y \leq y | O_1 = o_1, A_1 = a_1, O_2 = o_2, A_2 = a_2). \end{aligned}$$

This implies that the mean primary outcome of a DTR can be written as a function of the multivariate distribution of the observable data obtained from a SMART; see Murphy (2005a) for detailed derivation. This property ensures that data from SMARTs can be effectively used to evaluate pre-specified DTRs or to estimate the optimal DTR within a certain class. We defer our discussion of estimation of optimal DTRs to later chapters.

Power and Sample Size

As is the case with any other study, power and sample size calculations are crucial elements in designing a SMART. In a SMART, one can investigate multiple research questions, both concerning entire DTRs (e.g. comparing the effects of two DTRs) and concerning certain components thereof (e.g. testing the main effect of the first stage treatment, controlling for second stage treatment). To power a SMART, however, the investigator needs to choose a primary research question (primary hypothesis), and calculate the sample size based on that question. Additionally, one or more secondary questions (hypotheses) may be investigated in the study. While the SMART provides unbiased estimates (free from confounding) to these secondary questions by virtue of randomization, it is not necessarily powered to address these secondary hypotheses.

A good primary research question should be both scientifically important and helpful in developing a DTR. For example, in the addiction management study an interesting primary research question would be: "marginalizing over secondary treatments, what is the best initial treatment on average?". In other words, here the researcher wants to compare the mean primary outcome of the group of patients receiving NTX as the initial treatment with the mean primary outcome of those receiving CBT. Standard sample size formula for a large sample comparison of two means can be used in this case. Define the standardized effect size δ as the standardized difference in mean primary outcomes between two groups (Cohen 1988), i.e.

$$\delta = \frac{E(Y|A_1 = \text{NTX}) - E(Y|A_1 = \text{CBT})}{\sqrt{[Var(Y|A_1 = \text{NTX}) + Var(Y|A_1 = \text{CBT})]/2}}.$$

Suppose the randomization probability is $1/2$ for each treatment option at the first stage. Standard calculation yields a total sample size formula for the two sided test with power $(1-\beta)$ and size α:

$$n = 4(z_{\alpha/2} + z_\beta)^2 \delta^{-2},$$

where $z_{\alpha/2}$ and z_β are the standard normal $(1-\alpha/2)$ percentile and $(1-\beta)$ percentile, respectively. To use the formula, one needs to postulate the effect size δ, as is the case in standard two-group randomized controlled trials (RCTs).

Another interesting primary question could be: "on average what is the best secondary treatment, TM or TMC, for responders to initial treatment?". In other words, the researcher wants to compare the mean primary outcomes of two groups of responders (those who get TM versus TMC as the secondary treatment). As before, standard formula can be used. Define the standardized effect size δ as the standardized difference in mean primary outcomes between two groups (Cohen 1988), i.e.

$$\delta = \frac{E(Y|Response, A_2 = \text{TM}) - E(Y|Response, A_2 = \text{TMC})}{\sqrt{[Var(Y|Response, A_2 = \text{TM}) + Var(Y|Response, A_2 = \text{TMC})]/2}}.$$

Let γ denote the overall response rate to initial treatment. Suppose the randomization probability is $1/2$ for each treatment option at the second stage. Standard calculation yields a total sample size formula for the two sided test with power $(1-\beta)$ and size α:

$$n = 4(z_{\alpha/2} + z_\beta)^2 \delta^{-2} \gamma^{-1}.$$

To use the formula, one needs to postulate the overall initial response rate γ, in addition to postulating the effect size δ. A similar question could be a comparison of secondary treatments among non-responders; in this case the sample size formula would be a function of non-response rate to the initial treatment.

Alternatively researchers may be interested in primary research questions related to entire DTRs. In this case, Murphy (2005a) argued that the primary research questions should involve the comparison of two DTRs beginning with different initial treatments. Test statistics and sample size formulae for this type of research question have been derived by Murphy (2005a) and Oetting et al. (2011).

The comparison of two DTRs, say $\bar{d}$ and $\bar{d}'$, beginning with different initial treatments, can be obtained by comparing the subgroup of subjects in the trial whose treatment assignments are consistent with regime $\bar{d}$ with the subgroup of subjects in the trial whose treatment assignments are consistent with regime $\bar{d}'$. Note that there is no overlap between these two subgroups since a subject's initial treatment assignment can be consistent with only one of $\bar{d}$ or $\bar{d}'$. The standardized effect size in this context is defined as $\delta = (\mu_{\bar{d}} - \mu_{\bar{d}'})/\sqrt{(\sigma^2_{\bar{d}} + \sigma^2_{\bar{d}'})/2}$, where $\mu_{\bar{d}}$ is the mean primary outcome under the regime $\bar{d}$ and $\sigma^2_{\bar{d}}$ is its variance. Suppose the randomization probability for each treatment option is $1/2$ at each stage. In this case, using a large sample approximation, the required sample size for the two sided test with power $(1-\beta)$ and size α is

$$n = 8(z_{\alpha/2} + z_{\beta})^2 \delta^{-2}.$$

Oetting et al. (2011) discussed additional research questions and the corresponding test statistics and sample size formulae under different working assumptions. A web application that calculates the required sample size for sizing a study designed to discover the best DTR using a SMART design for continuous outcomes can be found at

`http://methodologymedia.psu.edu/smart/samplesize.`

Some alternative approaches to sample size calculations can be found in Dawson and Lavori (2010, 2012).

Furthermore, for time-to-event outcomes, sample size formulae can be found in Feng and Wahed (2009) and Li and Murphy (2011). A web application for sample size calculation in this case can be found at

`http://methodologymedia.psu.edu/logranktest/samplesize.`

Randomization Probabilities

Let $p_1(a_1|H_1)$ and $p_2(a_2|H_2)$ be the randomization probability at the first and second stage, respectively. Formulae for the randomization probabilities that would create equal sample sizes across all DTRs were derived by Murphy (2005a). This was motivated by the classical large sample comparison of means for which, given equal variances, the power of a test is maximized by equal sample sizes. Let $k_1(H_1)$ be the number of treatment options at the first stage with history H_1 and $k_2(H_2)$ be the number of treatment options at the second stage with history H_2, respectively. Then Murphy's calculations give the optimal values of randomization probabilities as

$$\begin{aligned} p_2(a_2|H_2) &= k_2(H_2)^{-1}, \text{ and} \\ p_1(a_1|H_1) &= \frac{E[k_2(H_2)^{-1}|O_1, A_1 = a_1]^{-1}}{\sum_{b=1}^{k_1(H_1)} E[k_2(H_2)^{-1}|O_1, A_1 = b]^{-1}}. \end{aligned} \tag{2.1}$$

If k_2 does not depend on H_2, the above formulae can be directly used at the start of the trial. Otherwise, working assumptions concerning the distribution of O_2 given (O_1, A_1) are needed in order to use the formulae. In the case of the addiction management example, $k_1(H_1) = 2$ and $k_2(H_2) = 2$ for all possible combinations of (H_1, H_2). Thus (2.1) yields an optimal randomization probability of $1/2$ for each treatment option at each stage. See Murphy (2005a) for derivations and further details.

2.3.3 Practical Considerations

Over the years, some principles and practical considerations have emerged mainly from the works of Lavori and Dawson (2004), Murphy (2005a) and Murphy et al. (2007a) which researchers should keep in mind as general guidelines when designing a SMART.

First, Murphy (2005a) recommended that the primary research question should consider simple DTRs, leading to tractable sample size calculations. For example, in the addiction management study, one can consider regimes where the initial decision rule does not depend on an individual's pre-treatment information and the secondary decision rule depends only on the individual's initial treatment and his response status (as opposed to depending on a large number of intermediate variables).

Second, when designing the trial, the class of treatment options at each stage should be restricted by ethical, scientific or feasibility considerations (Lavori and Dawson 2004; Murphy 2005a). It is better to use a low dimensional summary criterion (e.g. response status) instead of all intermediate outcomes (e.g. improvement of symptom severity, side-effects, adherence etc.) to restrict the class of possible treatments; in many contexts including mental health studies, feasibility considerations may often force researchers to use a patient's preference in this low dimensional summary. Lavori and Dawson (2004) demonstrated how to constrain treatment options (and thus decision rules) using the STAR*D study as an example (this study will be introduced later in this chapter). Yet, Murphy (2005a) warned against unnecessary restriction of the class of the decision rules. In our view, determining the "right class" of treatment options in any given study remains an art, and cannot be fully operationalized.

Third, a SMART should be viewed as one trial among a series of randomized trials intended to develop and/or refine a DTR (Collins et al. 2005). It should eventually be followed by a confirmatory randomized trial that compares the developed regime and an appropriate control (Murphy 2005a; Murphy et al. 2007a).

Fourth, like traditional randomized trials, SMARTs may involve usual problems such as dropout, non-compliance, incomplete assessments, etc. However, by virtue of the option to alter the non-functioning treatments at later stages, SMARTs should be more appealing to participants, which may result in greater recruitment success, greater compliance, and lower dropout compared to a standard RCT.

Finally, as in the context of any standard randomized trial, feasibility and acceptability considerations relating to a SMART can best be assessed via (external) *pilot studies* (see, e.g. Vogt 1993). Recently Almirall et al. (2012a) discussed how to effectively design a SMART pilot study that can precede, and thereby aid in fine-tuning, a full-blown SMART. They also presented a sample size calculation formula useful for designing a SMART pilot study.

2.3.4 SMART Versus Other Designs

The SMART design discussed above involves stages of treatment and/or experimentation. In this regard, it bears similarity with some other common designs, including what are known as *adaptive designs* (Berry 2001, 2004). Below we discuss the distinctions between SMART and some other multi-stage designs, to avoid any confusion.

SMART Design Versus Adaptive Designs

"Adaptive design" is an umbrella term used to denote a variety of trial designs that allow certain trial features to change from an initial specification based on accumulating data (evolving information) while maintaining statistical, scientific, and ethical integrity of the trial (Dragalin 2006; Chow and Chang 2008). Some common types of adaptive designs are as follows. A *response adaptive design* allows modification of the randomization schedules based on observed data at pre-set interim times in order to increase the probability of success for future subjects; Berry et al. (2001) discussed an example of this type of design. A *group sequential design* (Pocock 1977; Pampallona and Tsiatis 1994) allows premature stopping of a trial due to safety, futility and/or efficacy with options of additional adaptations based on the results of interim analyses. A *sample size re-estimation design* involves the re-calculation of sample size based on study parameters (e.g. revised effect size, conditional power, nuisance parameters) obtained from interim data; see Banerjee and Tsiatis (2006) for an example. An *adaptive dose-finding design* is used in early phase clinical development to identify the minimum effective dose and the maximum tolerable dose, which are then used to determine the dose level for the next phase clinical trials (see for example, Chen 2011). An *adaptive seamless phase II/III trial design* is a design that addresses within a single trial objectives that are normally achieved through separate trials in phase II and phase III of clinical development, by using data from patients enrolled before and after the adaptation in the final analysis; see Levin et al. (2011) for an example. In general, the aim of adaptive designs is to improve the quality, speed and efficiency of clinical development by modifying one or more aspects of a trial. Recent perspectives on adaptive designs can be found in Coffey et al. (2012).

Based on the above discussion, now we can identify the distinctions between the standard SMART design and adaptive designs. In a SMART design, each subject moves through multiple stages of treatment, while in most adaptive designs each stage involves different subjects. The goal of a SMART is to develop a good DTR that could benefit *future* patients. Many adaptive designs (e.g. response adaptive design) try to provide the most efficacious treatment to each patient *in the trial* based on the current knowledge available at the time that a subject is randomized. In a SMART, unlike in an adaptive design, the design elements such as the final sample size, randomization probabilities and treatment options are pre-specified. Thus, SMART designs involve *within-subject adaptation* of treatment, while adaptive designs involve *between-subject adaptation*.

Next comes the natural question of whether some adaptive features can be integrated into the SMART design framework. In some cases the answer is *yes*, at least in principle. For example, Thall et al. (2002) provided a statistical framework for an adaptive design in a multi-stage treatment setting involving two SMARTs. Thall and Wathen (2005) considered a similar but more flexible design where the randomization criteria for each subject at each stage depended on the data from all subjects previously enrolled. However, adaptation based on interim data is less feasible in settings where subjects' outcomes may only be observed after a long period of time has elapsed. How to optimally use adaptive design features within the SMART framework is an open question that warrants further research.

SMART Design Versus Crossover Trial Design

SMART designs have some operational similarity with classical crossover trial designs; however they are very different conceptually. First, treatment allocation at any stage after the initial stage of a SMART typically depends on a subject's intermediate outcome (response/non-response). However, in a crossover trial, subjects receive all the candidate treatments irrespective of their intermediate outcomes. Second, as the goal of a typical cross-over study is to determine the outcome of a one-off treatment, crossover trials consciously attempt to *wash out* the *carryover effects* (i.e. delayed effects), whereas SMARTs attempt to capture them and, where possible, take advantage of any interactions between treatments at different stages to optimize outcome following a sequence of treatments.

SMART Design Versus Multiphase Experimental Approach

As mentioned earlier, a SMART should be viewed as one trial among a series of randomized trials intended to develop and/or refine a DTR. It should eventually be followed by a confirmatory randomized trial that compares the developed regime and an appropriate control (Murphy 2005a; Murphy et al. 2007a). This purpose is shared by the *multiphase experimental approach* (with distinct phases for screening, refining, and confirming) involving factorial designs, originally developed in engineering

(Box et al. 1978), and recently used in the development of multicomponent behavioral interventions (Collins et al. 2005, 2009; Chakraborty et al. 2009). Note that DTRs are multicomponent treatments, and SMARTs are developmental trials to aid in the innovation of optimal DTRs. From this perspective, a SMART design can be viewed as one screening/refining experiment embedded in the entire multiphase experimental approach. In fact, Murphy and Bingham (2009) developed a framework to connect SMARTs with factorial designs. However, there remain many open questions in this context, and more research is needed to fully establish the connections.

2.4 Examples of Sequentially Randomized Studies

In this section, we consider two examples of SMARTs in great detail. An in-depth discussion of several other recently-conducted SMARTs can be found in Lei et al. (2012).

2.4.1 Project Quit – Forever Free: A Smoking Cessation Study

Here we briefly present a two-stage SMART design implemented in a study to develop/compare internet-based interventions (dynamic treatment regimes) for smoking cessation and relapse prevention. The study was conducted by the Center for Health Communications Research at the University of Michigan, and was funded by the National Cancer Institute (NCI). This study allowed the researchers to test cutting-edge web-based technology in a real-world environment that has the infrastructure for both evaluating and disseminating population-based cancer prevention and control programs. The first stage of this study, known as *Project Quit*, was conducted to find an optimal multi-factor behavioral intervention to help adult smokers quit smoking; and the second stage, known as *Forever Free*, was a follow-on study to help those (among the *Project Quit* participants) who had already quit remain non-smoking, and offer a second chance to those who failed to give up smoking at the previous stage. Details of the study design and primary analysis of the stage 1 data can be found in Strecher et al. (2008). Analysis of the data from the two stages considered together with a goal of finding an optimal DTR can be found in Chakraborty (2009) and Chakraborty et al. (2010).

At stage 1, although there were five two-level treatment factors in the original *fractional factorial* design, only two, `source` (of online behavioral counseling message) and `story` (of a hypothetical character who succeeded in quitting smoking) were significant in the primary analysis reported in Strecher et al. (2008). For simplicity of discussion, here we consider only these two treatment factors at stage 1, which would give a total of 4 treatment combinations at stage 1 corresponding to the 2×2 design. The treatment factor `source` was varied at two levels, e.g. high vs. low level of personalization; likewise the factor `story` was varied at two levels, e.g. high vs. low tailoring depth (degree to which the character in the story was tailored

to the individual subject's baseline characteristics). Baseline variables at this stage included subjects' `motivation` to quit (on a 1–10 scale), `selfefficacy` (on a 1–10 scale) and `education` (binary, ≤high school vs.> high school). At stage 2, there were two treatment options: booster intervention and control. At the first stage, 1,848 subjects were randomized, out of which only 479 decided to continue to stage 2 and hence were subsequently randomized.

There was an outcome measured at the end of each stage in this study. The stage 1 outcome was binary quit status at 6 months from the date of initial randomization, called `PQ6Quitstatus` (1 = quit, 0 = not quit). The stage 2 outcome was binary quit status, called `FF6Quitstatus`, at 6 months from the date of stage 2 randomization (i.e., 12 months from the date of stage 1 randomization). We will re-visit this study in Sects. 3.4.3 and 8.3.3, in the context of estimating optimal DTRs and conducting inference about them.

*2.4.2 STAR*D: A Study of Depression*

Sequenced Treatment Alternatives to Relieve Depression (STAR*D) was a multi-site, multi-level randomized controlled trial designed to assess the comparative effectiveness of different treatment regimes for patients with major depressive disorder (MDD) (Fava et al. 2003; Rush et al. 2004). This study was funded by the National Institute of Mental Health (NIMH). The study enrolled a total of 4,041 patients, all of whom were treated with citalopram (CIT) at level 1. Clinic visits occurred several times during each treatment level, at 2- or 3-week intervals (weeks 0, 2, 4, 6, 9, 12). Severity of depression at any clinic visit was assessed using the clinician-rated and self-report versions of the *Quick Inventory of Depressive Symptomatology* (QIDS) scores (Rush et al. 2004). A schematic of the treatment assignment algorithm is given in Fig. 2.3. This study is more complex than the smoking cessation study in that there are more than two stages.

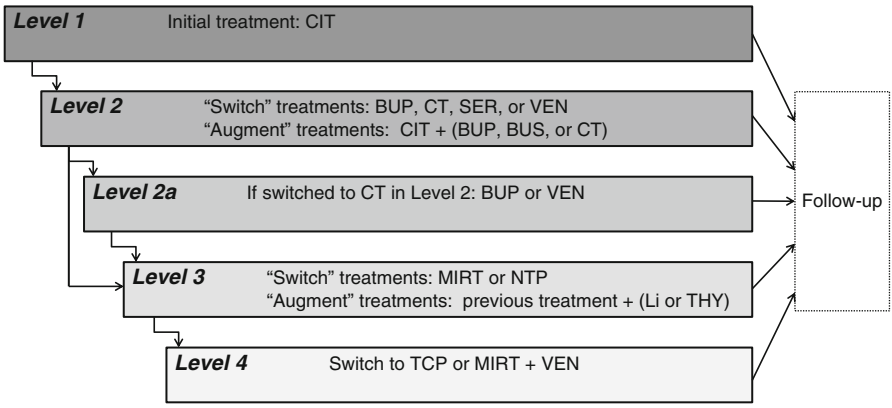

Fig. 2.3 A schematic of the algorithm for treatment assignment in the STAR*D study

Success was based on a total clinician-rated QIDS-score of ≤ 5 ("remission") during treatment with CIT. Those without remission were eligible to receive one of up to seven treatment options available at level 2, depending on their *preference* to switch or augment their level-1 treatment. Patients preferring a switch were randomly assigned to one of four treatment options: bupropion (BUP), cognitive psychotherapy (CT), sertraline (SER), or venlafaxine (VEN). Those preferring an augmentation were randomized to one of three options: CIT + BUP, CIT + buspirone (BUS), or CIT + CT. Only the patients assigned to CT or CIT + CT in level 2 were eligible, in the case of a non-satisfactory response, to move to a supplementary level of treatment (level 2A), to switch to either VEN or BUP. Patients not responding satisfactorily at level 2 (and level 2A, if applicable) would continue to level 3 treatment. Depending on the preference, patients at level 3 were randomly assigned to switch to either mirtazapine (MIRT) or nortriptyline (NTP), or randomly assigned to augment their previous treatment with lithium (Li) or thyroid hormone (THY). Patients without a satisfactory response at level 3 continued to level 4 treatments, which included two options: tranylcypromine (TCP) or MIRT + VEN. Patients achieving remission (QIDS ≤ 5) at any level entered a follow-up phase. Treatment assignment at each level took place via randomization within a patient's preference category. For a complete description of the STAR*D study design, see Fava et al. (2003) and Rush et al. (2004). We will re-visit this study in Chap. 8 in the context of making inference about the parameters indexing the optimal DTRs.

2.5 Discussion

In this chapter, we have described the two sources of data that are commonly used for estimating DTRs: observational follow-up studies and SMARTs. The use of observational data adds an element of complexity to the problem of estimation and requires careful handling and additional assumptions, due to the possibility of confounding. To assist in the careful formulation of causal contrasts in the presence of confounding, the potential outcomes framework was introduced. In contrast, SMARTs offer simpler analyses but often require significant investment to conduct a high quality trial with adequate power. We discussed conceptual underpinnings of and practical considerations for conducting a SMART, as well as its distinctions from other multiphase designs. We introduced several examples of observational and sequentially randomized studies, some of which we will investigate further in subsequent chapters.