

## Clarification: Model Specification and Estimation with Feasible Sets

In Chapters 6 and 7, we discuss considerations for the specification of outcome regression models and propensity models at each of the  $K$  decision points when interest focuses on a particular choice of feasible sets  $\Psi = (\Psi_1, \dots, \Psi_K)$  and for estimation of the value of a fixed regime and of an optimal regime in the class of  $\Psi$ -specific treatment regimes. As in Section 6.2.2, for  $k = 1, \dots, K$ ,  $\Psi_k(h_k) \subseteq \mathcal{A}_k$  comprises the options in  $\mathcal{A}_k$  that are feasible for an individual with history  $h_k$ . A feasible set  $\Psi_k(h_k)$  might contain all options in  $\mathcal{A}_k$  or a subset of the options in  $\mathcal{A}_k$ ; as in the ADHD example on page 218, the subset might comprise a single option in  $\mathcal{A}_k$ .

The discussions of modeling and estimation considerations in Chapters 6 and 7 are relevant under certain conditions regarding the nature of the observed data. Here, we present clarification of these conditions and describe how the developments should be modified if these conditions are not met. The **R** package **DynTxRegime** implements estimation of an optimal  $\Psi$ -specific regime correctly in either case and requires that the user understand the form of the observed data to specify implementation of the desired estimation method.

### Observed Data Scenarios

Assume interest focuses on a choice of feasible sets  $\Psi$ . As is discussed in Section 6.2.2, specification of  $\Psi$  is based ideally on scientific considerations. With this in mind, we distinguish two observed-data scenarios:

- (i) At each decision point  $k$ ,  $k = 1, \dots, K$ , individuals in the observed data with history  $h_k$  received only treatment options contained in  $\Psi_k(h_k)$ . This condition would be satisfied by default for observed data from a SMART focused on evaluation of  $\Psi$ -specific regimes in which subjects were randomized to the options in  $\Psi_k(h_k)$  based on their histories  $h_k$ . This condition would also be satisfied for observational data when the choice of  $\Psi$  reflects conventions in practice that dictate that patients with certain characteristics represented in their histories  $h_k$  would never receive treatment options not included in  $\Psi_k(h_k)$  at any decision point  $k$ ; e.g., if a treatment option would never be administered to an individual with certain characteristics on ethical or scientific grounds, as in the acute leukemia example, where salvage therapy would never be administered to a patient who had responded to induction therapy.

In general, this condition would be satisfied if  $\Psi$  were chosen to be the “largest” specification  $\Psi^{max}$  discussed at the end of Section 6.2.4. By definition,  $\Psi^{max}(h_k)$  includes all options received by individuals with history  $h_k$  in the observed data. When the available data are observational, it is common in practice to take  $\Psi = \Psi^{max}$ , so that  $\Psi_k(h_k)$  and  $\Psi^{max}(h_k)$  coincide for all  $k = 1, \dots, K$ .

- (ii) At at least one decision point  $k$ ,  $k = 1, \dots, K$ , there may be individuals in the observed data with history  $h_k$  who received treatment options *not contained* in  $\Psi_k(h_k)$ . This scenario is exemplified by the Remark on page 193 of Section 6.2.2 in the context of HIV treatment. Here,  $\Psi_k(h_k)$  is chosen based on the scientific premise that patients whose  $h_k$  indicates that the virus has become resistant to antiretroviral therapy should never continue to receive it because therapy is pointless and in fact detrimental when the virus is resistant. However, there may be patients with  $h_k$  indicating resistance

who insist on receiving therapy nonetheless.

More generally, there may be circumstances in which feasible sets may be deliberately chosen to restrict the feasible options for certain histories  $h_k$ ; e.g., reflecting interest in regimes to be used in resource-limited settings where certain options in  $\mathcal{A}_k$  would not be available. However, these options may have been administered to individuals with these histories in the observed data on which estimation is to be based.

Formally, under Scenario (ii), the specification of feasible sets  $\Psi$  chosen by the analyst is a strict subset of  $\Psi^{max}$ ; that is, for at least one  $k$ ,  $k = 1, \dots, K$ ,  $\Psi_k(h_k)$  is a strict subset of  $\Psi_k^{max}(h_k)$ ,  $\Psi_k(h_k) \subset \Psi_k^{max}(h_k)$ , for some  $h_k$ . Write this succinctly as  $\Psi \subset \Psi^{max}$ .

In Chapters 6 and 7, the discussions of model specification and estimation of a fixed  $\Psi$ -specific regime and of an optimal  $\Psi$ -specific regime tacitly assume that the observed data are consistent with Scenario (i). Thus, the presentation in each of the portions of the book indicated below is relevant under this condition but may not be relevant if the data are consistent with Scenario (ii).

## Modifications Under Scenario (ii)

When the observed data are consistent with Scenario (ii), the modeling and estimation considerations in the following sections of the book must be modified.

- (a) Section 6.4.2, implementation of the Backward Induction Approach to estimation of a fixed  $\Psi$ -specific regime
- (b) Section 6.4.3, Considerations for Propensity Modeling
- (c) Section 7.4.1, Modeling and Implementation Considerations for Q-learning
- (d) Section 7.4.2, Modeling and Implementation Considerations for A-learning
- (e) Section 7.4.3, Value Search Estimation

We now describe in detail how modeling and estimation (e.g., implemented using **DynTxRegime**) would proceed in practice under Scenario (ii). Suppose interest focuses on a specification of feasible sets  $\Psi \subset \Psi^{max}$ , and, for  $k = 1, \dots, K$ ,  $\mathcal{A}_k$  comprises  $m_k$  options. In what follows, we define as in the book  $\ell_k$  to be the number of distinct subsets  $\mathcal{A}_{k,l} \subseteq \mathcal{A}_k$ ,  $l = 1, \dots, \ell_k$ , that are feasible sets at Decision  $k$  under  $\Psi$ ;  $s_k(h_k)$  to take on values  $1, \dots, \ell_k$  according to which of these subsets  $\Psi_k(h_k)$  corresponds for given  $h_k$ , and  $M_k(h_k)$  to denote the number of options in  $\Psi_k(h_k)$ . The distinct subsets  $\mathcal{A}_{k,l}$  comprise  $m_{k,l}$  options,  $1 \leq m_{k,l} \leq m_k$ ,  $l = 1, \dots, \ell_k$ .

**Regression Modeling.** First consider the backward induction approach to estimation of the value of a fixed  $\Psi$ -specific regime in (a) and the method of Q-learning for estimation of an optimal  $\Psi$ -specific regime in (c). The discussion here focuses on Q-learning in (c), but the same considerations apply to (a).

As discussed in (c), a practical approach is to posit  $\ell_k$  separate models  $Q_{k,l}(h_k, a_k; \beta_{k,l})$ ,  $l = 1, \dots, \ell_k$ , for each subset, to arrive at an overall model as in (7.68) (equivalently (6.71) for (a)). With the parameters  $\beta_{k,l}$  sharing no common components across  $l = 1, \dots, \ell_k$  (i.e.,  $\beta_{k,l}$  are variationally independent), each

model can be fitted separately to the observed data from individuals with  $s_k(h_k) = l$  and used to form pseudo outcomes as discussed in (c).

At Decision  $k$ , suppose  $\Psi_k(h_k)$  for a given  $h_k$  contains  $M_k(h_k)$  options; equivalently, the distinct subset  $\mathcal{A}_{k,l}$  to which  $\Psi_k(h_k)$  corresponds has  $m_{k,l}$  options. However, in the observed data, there are individuals who received treatment options other than these  $M_k(h_k)/m_{k,l}$  options. We distinguish two cases:

- $M_k(h_k) > 1$ ; equivalently  $m_{k,l} > 1$ . Here, as above, the analyst would posit a regression model  $Q_{k,l}(h_k, a_k; \beta_{k,l})$ , where  $a_k$  can take on values corresponding to all treatments received in the observed data. The model would be fitted to the data on all individuals with  $s_k(h_k) = l$ , including the individuals receiving treatment options other than those in  $\Psi_k(h_k)$ . The fitted model would be used to obtain pseudo outcomes as described in (c).
- $M_k(h_k) = 1$ ; equivalently,  $m_{k,l} = 1$ . Under Scenario (i), as described in (c), no model would be posited and fitted, and the outcome ( $k = K$ ) or pseudo outcome  $\tilde{V}_{k,i}$  would be carried backward to Decision  $k - 1$ . Under Scenario (ii), this is inappropriate, as this outcome or pseudo outcome for individuals with this  $h_k$  who received a treatment option not in  $\Psi_k(h_k)$  will be inconsistent with having received the single option in  $\Psi_k(h_k)$ .

Accordingly, instead, the analyst should posit a regression model  $Q_{k,l}(h_k, a_k; \beta_{k,l})$ , where  $a_k$  can take on values corresponding to all treatments received in the observed data. The model would be fitted to the data on all individuals with  $s_k(h_k) = l$ , including the individuals receiving treatment options other than that in  $\Psi_k(h_k)$ . The fitted model would be used to obtain pseudo outcomes.

**Propensity Modeling.** Now consider propensity modeling used in estimation of the value of a fixed regime via an inverse probability weighted or augmented inverse probability weighted estimator as in (b) and estimation of an optimal  $\Psi$ -specific restricted regime using value search estimation based on either type of estimator in (e). The following considerations also apply to any method based on inverse weighting by propensity models.

As discussed in (b) and (e), a practical approach is to posit  $\ell_k$  separate propensity models  $\omega_{k,l}(h_k, a_k; \gamma_{k,l})$ ,  $l = 1, \dots, \ell_k$ , as in (6.105) and (7.93) for each subset to arrive at an overall model as in (6.106). The models  $\omega_{k,l}(h_k, a_k; \gamma_{k,l})$  may be logistic or multinomial (polytomous) logistic regression models, as discussed momentarily. With the parameters  $\gamma_{k,l}$  sharing no common components across  $l = 1, \dots, \ell_k$  (i.e.,  $\gamma_{k,l}$  are variationally independent), each model can be fitted separately by maximum likelihood to the observed data from individuals with  $s_k(h_k) = l$ .

At Decision  $k$ , suppose  $\Psi_k(h_k)$  for a given  $h_k$  contains  $M_k(h_k)$  options; equivalently, the distinct subset  $\mathcal{A}_{k,l}$  to which  $\Psi_k(h_k)$  corresponds has  $m_{k,l}$  options. However, in the observed data, there are individuals who received treatment options other than these  $M_k(h_k)/m_{k,l}$  options. We distinguish three cases:

- $M_k(h_k) = 2$ ; equivalently,  $m_{k,l} = 2$ . Under Scenario (i), as above, one would posit and fit a logistic regression model  $\omega_{k,l}(h_k, a_k; \gamma_{k,l})$  for the  $M_k(h_k) = 2$  treatment options. If, as here, there are individuals who received one or more options not contained in  $\Psi_k(h_k)$ , one would instead take  $\omega_{k,l}(h_k, a_k; \gamma_{k,l})$  to be a multinomial (polytomous) logistic regression model for the total number of treatment options ( $> 2$ ) received by individuals with  $s_k(h_k) = l$  in the observed data, including the options not in  $\Psi_k(h_k)$ . The model would be fitted to the data on all of the individuals with  $s_k(h_k) = l$ .

- $M_k(h_k) > 2$ ; equivalently,  $m_{k,l} > 2$ . Under Scenario (i), one would posit and fit a multinomial (polytomous) logistic regression  $\omega_{k,l}(h_k, a_k; \gamma_{k,l})$  for the  $M_k(h_k) > 2$  treatment options. If, as here, there are individuals who received one or more options not contained in  $\Psi_k(h_k)$ , one would instead take  $\omega_{k,l}(h_k, a_k; \gamma_{k,l})$  to be a multinomial (polytomous) logistic regression model for the total number of treatment options ( $> M_k(h_k)$ ) received by individuals with  $s_k(h_k) = l$  in the observed data, including the options not in  $\Psi_k(h_k)$ . The model would be fitted to the data on all of the individuals with  $s_k(h_k) = l$ .
- $M_k(h_k) = 1$ ; equivalently,  $m_{k,l} = 1$ . Under Scenario (i), as described in (b) and (e), the propensity for the single option in  $\Psi_k(h_k)$  is set equal to 1. Under Scenario (ii), this is inappropriate, as the probability of receiving this single option under the observed data treatment assignment mechanism is  $< 1$ .

Accordingly, the analyst should posit and fit a logistic or multinomial (polytomous) logistic regression model  $\omega_{k,l}(h_k, a_k; \gamma_{k,l})$  as appropriate for the total number of treatment options ( $> 1$ ) received by individuals with  $s_k(h_k) = l$  in the observed data, including the options not in  $\Psi_k(h_k)$ . The model would be fitted to the data on all of the individuals with  $s_k(h_k) = l$ .

For A-learning in (d), under Scenario (ii), the propensity models should be developed according to these considerations. In addition, under Scenario (ii), the contrast function models should be developed in a manner analogous to that for Q-function models in (a) and (c) above. In particular, if  $M_k(h_k) = 1$  or equivalently  $m_{k,l} = 1$ , but there are individuals in the observed data with  $s_k(h_k) = l$  who received treatment options different from the one in  $\Psi(h_k)/\mathcal{A}_{k,l}$ , a contrast function model should be posited and fitted based on the data from all such individuals and used to construct pseudo outcomes rather than “carrying backward” pseudo outcomes as described in (d) under Scenario (i).

## Implementation in Practice

The foregoing considerations dictate that it is incumbent on a data analyst interested in a particular choice of feasible sets  $\Psi = (\Psi_1, \dots, \Psi_K)$  and the class of  $\Psi$ -specific treatment regimes to scrutinize the observed data and determine which of Scenario (i) or (ii) is relevant. This exercise will inform the analyst in advance of the need to posit suitable models to be fitted to the data on subjects whose histories are consistent with each feasible set.

In **DynTxRegime**, under Scenario (i), the analyst should specify models for each feasible set with at least two treatment options. Under Scenario (ii), a model should be specified for each feasible set for which individuals in the data whose histories are consistent with the feasible set are observed to have received more than one option, even if the feasible set itself comprises only one option. Under either scenario, the models should be specified according to the to the guidelines given above.

**DynTxRegime** makes every effort to identify possible issues between the models specified for an analysis and the data provided. Specifically, for outcome regression modeling:

- (i) For  $M_k(h_k) > 1$ , if individuals with  $s_k(h_k) = l$  received treatment options other than those in  $\Psi_k(h_k)$  a message is generated. This message is informative only and does not stop the analysis.

- (ii) For  $M_k(h_k) = 1$ , if individuals with  $s_k(h_k) = l$  received treatment options other than those in  $\Psi_k(h_k)$  execution is halted with an error message indicating that a model must be provided for individuals with  $s_k(h_k) = l$ .

And for propensity regression modeling:

- (i) For  $M_k(h_k) \geq 2$ , if individuals with  $s_k(h_k) = l$  received treatment options other than those in  $\Psi_k(h_k)$  a message is generated. This message is informative only and does not stop the analysis. Note that the modeling object framework makes it difficult if not impossible to determine in general if a model is a logistic regression or a multinomial regression. Therefore, the scenario where a logistic regression is provided but a multinomial is required does not result in an information proactive stop but will result in a message and likely a regression error.
- (ii) For  $M_k(h_k) = 1$ , if individuals with  $s_k(h_k) = l$  received treatment options other than those in  $\Psi_k(h_k)$  execution is halted with an error message indicating that a model must be provided for individuals with  $s_k(h_k) = l$ .