Estimating dynamic treatment regimes.

Matthew Blackwell | Harvard University | mblackwell@iq.harvard.edu

The Problem

TREATMENTS ≠ SINGLE-SHOT

Causes of effects in political science are hardly ever single-shot. Treatments are set, re-set, and adjusted. Currently, there is no infrastructure for political scientists to handle this type of situation. I introduce a methodology based on Marginal Structural Models (MSMs) to handle dynamic treatments in political science. For example, in campaigns, campaign advertising tone is a dynamic treatment:

$$\begin{tikzcd} \text{Tone}_1 \arrow[r] \arrow[rrr, bend right] & \text{Polls}_1 \arrow[r] \arrow[rr, bend left] & \text{Tone}_2 \arrow[r] & \text{Votes} \end{tikzcd}$$

How would we control for polls in this model? Standard statistical advice would be to control for the polls because they confound the relationship between advertising and electoral outcomes. On the other hand, we should omit polls, because they are a consequence of advertising. Neither of these two approaches will consistently estimate causal effects. Moreover, estimates from both will fail to bound the effect in very typical political science environments.

The Concepts

Teatment Regimes

A rule, r, for assigning treatment in each round, possibly conditional on past covariates and treatment. Treatment regimes are very similar to strategies in dynamic game theory models. Examples include:

- Fixed: Always run negative ads.
- Static: Go negative early in the race.
- Triggers: Only go negative if my polls fall below x%.
- Tit-for-tat: Only go negative if attacked by opponent.

Potential Outcomes

A potential outcome, $Y_i(r)$, is the outcome that would have happened to unit i if they had followed regime r. Interestingly, units may have followed more than one regime. For instance, if a candidate went negative when polls went from 45% in their favor to 40%, then they would have followed the "trigger" regimes from above with $x = (41, \ldots, 45)$.

Notation

- A_t = treatment in time t
- $\bar{A}_t$ = treatment history up to time t
- X_t = covariates in time t
- $\bar{X}_t$ = covariate history up to time t

The Model

An MSM is a model for the mean of the potential outcomes as a function of the treatment regimes:

$$\mathbb{E}[Y_i(r)] = g(r; \beta).$$

This function might be a simple cumulative exposure for static regimes $r = \bar{a}$:

$$g(\bar{a}; \beta) = \beta_0 + \beta_1 \left(\sum_t a_t \right).$$

For dynamic regimes, it might compare two different regimes:

$$g(r; \beta) = \beta_0 + \beta_1 \mathbb{I}_r$$

$$\mathbb{I}_r = \begin{cases} 1 \text{ if } r = \text{tit-for-tat} \\ 0 \text{ if } r = \text{always positive} \end{cases}$$

This model is semiparametric in that we only restrict the mean of the potential outcomes, not the entire distribution.

The Estimation

Running the above simple regressions on the observed outcome Y_i will generally lead to incorrect estimates of the causal parameter β_1 whether or not you further adjust for time-varying confounders. Instead, we use an IPTW estimator with weights:

$$W_i = \frac{1}{\prod_t p(A_t | \bar{A}_{t-t}, \bar{X}_t; \hat{\alpha})},$$

where $\hat{\alpha}$ are logit-estimated parameters. The denominator is the estimated probability of receiving the treatment history that unit i actually did receive.

Procedure

1. Fit a pooled logit of current treatment A_t on treatment and covariate history to estimate α.
2. Remove any unit that did not follow any of the regimes being tested.
3. Fit the MSM with each observation weighted by W_i to estimate β:
$$\mathbb{E}[Y_i | r_i] = g(r_i; \beta)$$
4. Use a nonparametric bootstrap for estimates of uncertainty.

The Assumptions

Assumption 1 (Consistency). *For any treatment regime, observed outcomes are equal to the potential outcome under the treatment regime actually observed. Formally, if unit i has a treatment history consistent with regime r, then $Y_i = Y_i(r)$.*

Assumption 2 (Sequential Ignorability). *For any treatment regime r, time t, treatment assignment is independent of the potential outcome conditional on observed information available at t. Formally,*

$$Y(r) \perp\!\!\!\perp A_t | (A_{t-1}, \ldots, A_1), (X_{t-1}, \ldots, X_1), \forall t.$$

These extend SUTVA and ignorability to the time-varying context.

The Data

The data currently include 68 Senate and Governor elections in 2000 and 2002. Variables used include:

- Outcome: the democratic vote-share in the election.
- Data on every television advertisement shown by both candidates, including tone.
- Every publicly available poll for each race.
- Background covariates on each race and candidate.
- Predicted closeness of the race (CQ rating).

I used a subset of these variables to construct the weights, eliminating variables that seemingly had no effect on treatment. Passed intuition checks as negativity increases with: past negativity from both sides, the closeness of the race, the week of the campaign, and the total advertising volume.

The Results

I tested the effect of two types of strategies on the Democratic percent of the two-party vote.

- Effect of an additional week of negativity.
- Effect of "hitting back" within three weeks.

With both, I compared the MSM estimate to estimates from omitting and adjusting for confounders. The MSM finds greater effects than either of the two traditional methods.

Treatment Effects

	Adjust (SE)	Omit (SE)	MSM (SE)
Another week	0.37 (0.21)	0.87 (0.25)	0.91 (0.25)
Hitting back	4.9 (3.1)	7.6 (3.4)	9.1 (3.1)