

Copyright  
by  
Byung-Jae Lee  
2014

The Dissertation Committee for Byung-Jae Lee  
certifies that this is the approved version of the following dissertation:

**Methodological Problems in Causal Inference, With  
Reference to Transitional Justice**

Committee:

---

Robert C. Luskin, Supervisor

---

Daniel M. Brinks

---

Zachary S. Elkins

---

Tse-min Lin

---

Robert G. Moser

---

Daniel A. Powers

**Methodological Problems in Causal Inference, With  
Reference to Transitional Justice**

by

**Byung-Jae Lee, B.A.; M.A.; M.A.**

**DISSERTATION**

Presented to the Faculty of the Graduate School of  
The University of Texas at Austin  
in Partial Fulfillment  
of the Requirements  
for the Degree of

**DOCTOR OF PHILOSOPHY**

THE UNIVERSITY OF TEXAS AT AUSTIN

August 2014

Dedicated to my wife, my parents, my father-in-law and my late mother-in-law

# Acknowledgments

I would like to express my deepest gratitude to my dissertation committee members — Robert Luskin, Daniel Brinks, Zachary Elkins, Tse-min Lin, Robert Moser, and Daniel Powers — and David Leal for helping me in various stages of my life at Austin.

# Methodological Problems in Causal Inference, With Reference to Transitional Justice

Publication No. \_\_\_\_\_

Byung-Jae Lee, Ph.D.

The University of Texas at Austin, 2014

Supervisor: Robert C. Luskin

This dissertation addresses methodological problems in causal inference in the presence of time-varying confounding, and provides methodological tools to handle the problems within the potential outcomes framework of causal inference. The time-varying confounding is common in longitudinal observational studies, in which the covariates and treatments are interacting and changing over time in response to the intermediate outcomes and changing circumstances. The existing approaches in causal inference are mostly focused on static single-shot decision-making settings, and have limitations in estimating the effects of long-term treatments on the chronic problems. In this dissertation, I attempt to conceptualize the causal inference in this situation as a sequential decision problem, using the conceptual tools developed in decision theory, dynamic treatment regimes, and machine learning. I also provide methodological tools useful for this situation, especially when the treat-

ments are multi-level and changing over time, using inverse probability weights and  $g$ -estimation.

Substantively, this dissertation examines transitional justice's effects on human rights and democracy in emerging democracies. Using transitional justice as an example to illustrate the proposed methods, I conceptualize the adoption of transitional justice by a new government as a sequential decision-making process, and empirically examine the comparative effectiveness of transitional justice measures — independently or in combination with others — on human rights and democracy.

# Table of Contents

<b>Acknowledgments</b>	<b>v</b>
<b>Abstract</b>	<b>vi</b>
<b>List of Tables</b>	<b>xi</b>
<b>List of Figures</b>	<b>xii</b>
<b>Chapter 1. Introduction</b>	<b>1</b>
Organization of the Dissertation . . . . .	6
<b>Chapter 2. An Overview of Causal Inference and Its Decision-Theoretic Elements</b>	<b>8</b>
Potential Outcomes Framework . . . . .	9
Identification of Causal Effects . . . . .	14
Observational Studies . . . . .	15
Graphical Criteria for Identification . . . . .	16
Interpreting Causal Inference from Decision-Theoretic Perspective . . . . .	20
Subject-specific Treatment and Decision Making . . . . .	21
Decision in Single-stage . . . . .	21
Decisions in Multi-stage . . . . .	23
Methodological Implications for Longitudinal Studies . . . . .	24
Time-varying Confounding and Mediation . . . . .	25
Assumptions for Causal Inference in Sequential Setting . . . . .	27
Discussion . . . . .	31



<b>Chapter 3. Causal Effects of Transitional Justice</b>	<b>33</b>
Definitions of Transitional Justice . . . . .	33
Varieties of Transitional Justice Mechanisms . . . . .	35
Spread of Transitional Justice? . . . . .	41
Mechanisms of Transitional Justice Effects . . . . .	51
Temporal Dimensions of Transitional Justice Effects . . . . .	52
Discussion . . . . .	56
<b>Chapter 4. Using Propensity Score and Inverse Probability Weights</b>	<b>58</b>
Inverse Probability Weighted Estimators: Treatment and Censoring . . . . .	60
Propensity Scores and Inverse Probability Treatment Weights for Multiple Treatments . . . . .	64
The Effect of Human Rights Trials on Human Rights . . . . .	68
Data . . . . .	69
Analysis 1: Moving Beyond Binary Treatment . . . . .	74
Analysis 2: Moving to the Inclusion of Covariates and Adjustment . . . . .	78
Discussion . . . . .	85
<b>Chapter 5. Causal Inference for Dynamic Treatment Data</b>	<b>87</b>
The Methodological Framework to Estimate Dynamic Treatments . . . . .	90
Set-up and Notations . . . . .	90
Main Assumptions . . . . .	92
Modeling Dynamic Treatments with Marginal Modeling . . . . .	94
Marginal Structural Models and Inverse Probability Weighting . . . . .	94
Structural Nested Mean Model (SNMM) and $g$ -estimation . . . . .	96
Transitional Justice as Dynamic Treatment Regime . . . . .	101
Application of Inverse Probability Weights for Adjustment: Physical Integrity . . . . .	102
Application of $G$ -estimation: Democratic Regime Survival . . . . .	105
Discussion . . . . .	108

<b>Chapter 6. Dynamic Treatment Regime From Machine Learning Perspective</b>	<b>110</b>
Reinforcement Learning: A Brief Introduction . . . . .	112
Probabilistic Framework for Reinforcement Learning . . . . .	114
Estimating Optimal Dynamic Treatment Regime . . . . .	120
Simulation . . . . .	125
Application: Economy and Human Rights Trials . . . . .	132
Discussion . . . . .	134
<b>Chapter 7. Conclusion</b>	<b>136</b>
Implications . . . . .	139
Future Research . . . . .	141
<b>Bibliography</b>	<b>143</b>

## List of Tables

3.1	The Number of Countries that Adopted Each Transitional Justice Measure . . . . .	42
3.2	The Number of Countries in Each Combination of Transitional Justice Measures . . . . .	43
3.3	Years of Transitional Justice . . . . .	46
3.4	The Number of Years that Each Country Adopted Transitional Justice Measures . . . . .	46
4.1	The Number of Regimes and Democratic Regimes . . . . .	72
4.2	Physical Integrity Before and After Transitional Justice: 3, 5, and 10 years after Democratic Transition . . . . .	77
4.3	Physical Integrity Before and After Transitional Justice: Combined Measures . . . . .	78
4.4	Ordered Regression Result 1 . . . . .	83
4.5	Predicted Probabilities for Select Scenario (unadjusted) . . . . .	84
4.6	Predicted Probabilities for Select Scenario (adjusted) . . . . .	85
5.1	Ordinal Regression Result 2 . . . . .	104
5.2	Functional Forms of $g$ -estimation . . . . .	106
5.3	$G$ -estimation Results: Static Regime . . . . .	107
5.4	$G$ -estimation Results: Dynamic Regime . . . . .	108
6.1	Performance of Four Adjustment Methods . . . . .	129
6.2	Performance of Adjustment Methods under Confounding . . . . .	132
6.3	Coefficients in Two Stage Regression for $Q$ -learning . . . . .	133
6.4	Estimates and Confidence Intervals for Contrasts in Two Stages . . . . .	133

## List of Figures

2.1	A Two-stage DAG illustrating Time-varying Confounding and Mediating . . . . .	27
3.1	Choice for New Government on Transitional Justice . . . . .	38
3.2	The Number of Countries that Adopted the Transitional Justice Year for the Particular Year (beginning and continuing), 1970-2008 . . . .	44
3.3	The Cumulative Number of Countries that Take the Transitional Justice, 1970-2008 . . . . .	45
3.4	The Comparison of Density Plots of the Key Variables . . . . .	50
3.5	Timeline for the Ideal Progression of Transitional Justice . . . . .	55
3.6	Adopted Transitional Justice Measures according to the Regime Duration . . . . .	56
4.1	Single and Multiple Transitions . . . . .	70
4.2	Density Plots of Stabilized Weights for Each Treatment Level . . . .	81
4.3	Density Plots for Stabilized, Censoring, and Stabilized Censoring Weights	82
4.4	Predicted Probabilities . . . . .	85
5.1	The Chronological Order of the Data in Two Interval Model . . . . .	91
5.2	Directed Acyclic Graphs Representing Different Assumptions about Sequential Ignorability . . . . .	93
5.3	Density Plots for the Weights . . . . .	103
5.4	Predicted Probabilities according to the Years in Trial . . . . .	105
5.5	Kaplan-Meier Estimates for Regime Survival with Dynamic Regime .	108
6.1	Reinforcement Learning in Markov Decision Process . . . . .	113
6.2	Directed Acyclic Graph for Simulated Data . . . . .	127

# Chapter 1

## Introduction

This dissertation consists of multiple essays on causal inference and transitional justice. Although each essay attempts to deal with different problems in causal inference, the overarching theme is how to deal with time-varying or time-dependent confounding both from covariates and treatments.

Causal inference is a central goal of most, if not all, science. Social scientists in various disciplines, regardless of their methodological orientation, seek to uncover causal relationships underlying the social phenomena they are interested in. However, the meaning of something being “causal” is often not made explicit. What does it actually mean that  $x$  causes  $y$ ?

A typical causal hypothesis takes the following form: “ $x$  is a cause of  $y$ ,” or equivalently, “ $x$  causes  $y$ .” For example, one of the perennial questions in comparative politics is the causal relationship among democracy, economic development and political culture, and the questions have taken the form of whether economic development causes democracy or whether political culture has causal power for democracy (Inglehart and Welzel 2005; Przeworski et al. 2000). Although causality can be defined in a multitude of otherwise useful ways (Hoover 2001; Ehring 1997; Spirtes, Glymour and Scheines 2000; Beebe, Hitchcock and Menzies 2012), the coun-

terfactual conception of causality paved the way for the revived interests in causal inference in social science (Lewis 1973; Brady 2008; Sekhon 2004, 2008; Woodward 2003, 2009). According to this conception,  $x$  is a cause of  $y$  is equivalent to “ $y$  would not have happened in the absence of  $x$ .” Symbolically,  $x \rightarrow y$  can be interpreted counterfactually as “ $\neg x \rightarrow \neg y$ ” (No bourgeoisie, no democracy). Although the crucial importance lies in theoretical mechanism between  $x$  and  $y$ , the casual inference literature usually treats them as a black box (Glynn and Quinn 2011; Bullock, Green and Ha 2010; Gelman and Hill 2006).

Although very few people paid attention to the similarity between them, causal statements can be translated into a decision theoretic framework. In this dissertation, I attempt to combine causal inference with a decision-theoretic perspective, which opens the way to link with the dynamic treatment regime and reinforcement learning, two methods recently developed in biostatistics and computer science, respectively. If  $x$  causes  $y$  and  $y$  is the desired goal for the decision-maker, it automatically follows that the decision-maker attempts to maximize  $x$  by contrasting the outcomes. This decision rule becomes more complicated in observational studies, where treatment is not randomized and the decisions are made sequentially in response to the intermediate outcomes and the changing milieu. Although many treatment evaluations and game theory are based on static settings, no reasonable policy maker sticks to the policy whose intermediate performance is dismal! The changing intermediate decision in the middle poses serious problems in causal inference because causal inference is typically framed in a single shot setting, and the quantity of interest is treatment effects estimated by comparing some measures

before and after treatments. The problem is that the existing literature in causal inference is not well equipped to deal with this problem. The approach is to consider the treatments not as a single treatment, but as a sequence of treatments. In this framework, short-term loss (gain) might result in long-term gain (loss), and local maxima are not the global maxima. The effectiveness of a treatment sequence needs to be estimated at the end of the study, not at the end of each stage.

The methodological goal of this dissertation is to provide tools to estimate causal effects in presence of time-varying confounding. This attempt has significant implications for dealing with time in social science. The typical causal questions in political science take the form of “does economic development bring democracy?” (Lipset 1959; Przeworski et al. 2000) or “are consumer prices higher in majoritarian electoral systems than in proportional representation systems?” (Rogowski and Kayser 2002; Chang et al. 2010) Even in studies that deal with longitudinal or time-series data, time-varying confounding is not well-addressed. Although many economists and political scientists have paid attention to the effect of the timing and sequence of decisions on a country’s trajectory (Pierson 2004), most have focused on particular cause of certain events or institutions with small number of cases, and very few attempts have been made to address the timing and sequence quantitatively (Page 2006; Jackson and Kollman 2010).

Throughout the dissertation, I use transitional justice as an illustrative example. Specifically, I draw examples from the causal questions of whether and which transitional justice mechanisms affect human rights and the stability of democratic regimes.

A few introductory words, therefore, about “transitional justice.” The “transitions” involved are from periods of repression, civil war, genocide, or other intense protracted internal conflicts. Over the past several decades there have been more than a hundred such transitions,<sup>1</sup> characteristically revolving around widespread unhappiness with the outgoing regime’s transgressions, which both help spark the change of regime and remain a central issue in negotiating an eventual peace accord, formal pact,<sup>2</sup> or revised constitution. The new government, often but not always democratic, must inevitably decide what, if anything, to do about the accumulated grievances and residual memories of the previous one, in particular whether to hold the architects and perpetrators of past abuses accountable. Survivors, families of victims, repressed communities, and concerned members of civil society ardently call for justice and reparation. Yet the new government may face practical constraints, and approaches vary.

Transitional justice is hardly a new phenomenon,<sup>3</sup> but it was relatively recent, after the ‘third’ wave of democratization, that a serious discussion of transitional justice and its effects began. Given that any transition requires the new regime to deal with the past and that most of historical transitions include such typical components of transitional justice as purge, trial, and reparation, an important question is then

---

<sup>1</sup>Those processes involving democratization in spatio-temporal clusters have been described as a worldwide ‘third wave’ and ‘fourth wave’ of democratization (Huntington 1991; McFaul 2002).

<sup>2</sup>‘Pacted’ or ‘negotiated’ transition (or reform by transaction) is generally characterized by a negotiated compromise between the elites of the authoritarian regime and the democratic opposition (Munck and Leff 1997).

<sup>3</sup>Notable historical examples include ancient Athens (411 and 403 B.C.) and French Restoration in 1814 and 1815 (see Moore 1975; Elster 2004), and there were plenty of historical and legal discussions on transitional justice after the Second World War.



how recent transitional justice is different from that in the past, and whether and how the differences are conducive to democratic consolidation.

The increased diversity of transitional justice measure has significantly expanded the options for the new governments. It is only after the ‘third’ wave of democratization that the transitional justice was discussed as institutional<sup>4</sup> mechanism to come to terms with the past, especially after the introduction of truth commission as an alternative or supplementary mechanism to human rights trials or virtual immunity (inaction). The commonly used transitional justice measures now include (domestic, international and hybrid) human rights trials, truth commissions, reparations, rehabilitations, file access, vetting (lustration<sup>5</sup>), constitutional reform, reform in the security sector, implementation of ombudsman, public apology, memorials, museum, textbook reform, street naming, national holidays, among others, many of which were unavailable before 1980s.

Nevertheless, our knowledge on the effects of transitional justice is very limited despite the accumulated qualitative and quantitative data on transitional justice over the last two decades. Most of the previous studies on transitional justice were and still are ‘faith-based’ rather than ‘fact-based’ (Thoms, Ron and Paris 2010), and there are no solid theoretical framework and sufficient empirical evidence yet to

---

<sup>4</sup>I am using the term institution loosely to distinguish it from transitional justice as episodic events, because transitional justice cannot be considered an institution in the commonly used sense of the rule of the game governing the behavior of the actors (North 1990) or equilibrium behavior (Schotter 1981; Bates et al. 1998).

<sup>5</sup>A process of ‘purification’ that excludes various types of officials, functionaries and elites based on their actual or presumed complicity in past abuses from participation in the successor government or in the civil services.

examine the causal relationship between transitional justice and other democratic goals (Mendeloff 2004; Thoms, Ron and Paris 2010).

## **Organization of the Dissertation**

This dissertation consists of seven chapters. Chapter 2 is an overview of causal inference with specific focus on its relationship with decision theory. This chapter describes the assumptions, estimations, identification, and limitations of causal inference. It also briefly describes the time-varying confounding and mediation, and it argues that causal inference can be interpreted as a decision problem.

Chapter 3 is an overview of transitional justice, which discusses the definition and the patterns of transitional justice and the expected effects.

Chapter 4 is a discussion of propensity score and inverse probability treatment weight and their use in estimating the causal effects when selection effects need to be adjusted. I argue that propensity scores and inverse probability treatment weights can be usefully implemented as weights to adjust the bias due to selection. In this chapter, I extend the use of inverse probability weights from binary to multiple treatments, and apply to the estimation of transitional justice's effects on human rights.

Chapter 5 discusses dynamic treatment regimes, useful in modeling treatments for chronic problems requiring adaptive treatments. A decision-maker changes treatment strategy in response to the changing situations and intermediate outcomes. He or she may continue, stop, or adjust the ongoing treatment, and this continuing process needs to be understood as a set of treatments (a treatment regime) rather

than a single treatment. The estimation of the causal effect of whole series needs a new methodological strategy. I estimate the effects of ordinal multiple treatments of transitional justice sequence using inverse probability weights and  $g$ -estimation.

Chapter 6 is an attempt to combine dynamic treatment regime with reinforcement learning, a relatively less known branch of machine learning, with specific focus on causal inference. I argue that dynamic treatment regime and reinforcement learning are similar, and optimal dynamic treatment regime can easily be reformulated by reinforcement learning.

Chapter 7 is a conclusion, offering final thoughts on the implications and possible extensions of the analyses of the preceding chapters.

## Chapter 2

# An Overview of Causal Inference and Its Decision-Theoretic Elements

Suppose that a country maintains democratic stability after holding free and fair elections (Lindberg 2009). Did the voting cause the stability? According to the counterfactual conception of causality, a natural way to think about this question would be to imagine what would have happened had the elections not been held. If the country would have maintained democratic stability anyway, we would not say that the voting was the cause. If, on the other hand, the country would not have maintained the stability without the voting, then we would say that the voting caused the stability. Here, stability is the outcome, and the election is the treatment or action.<sup>1</sup> To determine whether a treatment or an action causes an outcome, we typically make a mental comparison between the two scenarios; one where the treatment is present and one where the treatment is absent. If the outcome differs between the two scenarios, we say that the treatment has a causal effect on the outcome. The potential outcomes framework formalizes this intuition of causality.<sup>2</sup>

---

<sup>1</sup>Treatment, action, and exposure are interchangeably used in causal inference literature, although epidemiologists prefer exposure, social scientists treatment.

<sup>2</sup>Counterfactual thought experiments have a long tradition in social science dating back at least to Max Weber for its systematic treatment (Weber 1949; Elster 1978; Fearon 1991; Tetlock and Belkin 1996). For a thorough discussion of requirement for meaningful counterfactual statement, see Tetlock and Belkin (1996).

## Potential Outcomes Framework

The potential outcomes framework for causal inference is based on a specific conception of causality, called counterfactual conception (Rubin 2006; Brady 2008; Sekhon 2008), and most of recent literature on causal inference relies on the notion of potential outcome, defined as an outcome had the subject followed a particular treatment, possibly different from the treatment he or she actually followed.

In experimental or clinical settings, the individual-level causal effect of a treatment may be viewed as the difference in outcomes if a person had followed that treatment as compared to a placebo or a standard protocol (Morton and Williams 2010). Consider, for example, a simple randomized trial in which subjects can receive either treatment  $a$  or  $a'$ . Suppose further that an individual was randomized to receive treatment  $a$ . This individual will have a single observed outcome  $Y$  that 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  $a'$ .

The so-called *fundamental problem of causal inference* (Holland 1986) lies in the definition of causal parameters at an individual level. Suppose we are interested in the causal effects of taking treatment  $a$  instead of treatment  $a'$ . The individual level causal parameter that could be considered is a subject's outcome under treatment  $a'$  subtracted from his outcome under treatment  $a$ , i.e.,  $Y(a) - Y(a')$  (subject-specific causal effect). If, for a given subject, all potential outcomes are equal (i.e.,  $Y$  does not depend on  $a$ ), then for this subject, the treatment has no causal effect on the outcome. If the treatment has no causal effect on the outcome for any subject in the study population, we could say the causal null hypothesis holds. A fundamental problem

with subject-specific causal effects is that they are difficult to identify, because it is difficult to observe the outcome under both  $a$  and  $a'$  without further data and assumptions, as in crossover designs without carryover effects (Piantadosi 2005, 515).

To generalize, let  $A$ ,  $Y$  and  $X$  denote the observed treatments, outcome, and covariates, respectively, for a given subject.

**Assumption 1.** *If a subject is treated to level  $A = a'$ , the potential outcome  $Y_{a'}$  is assumed to be equal to the observed factual outcome  $Y$  for that subject. This is called consistency assumption, and can be formally expressed:*

$$A = a' \Rightarrow Y_{a'} = Y \tag{2.1}$$

We remain ignorant, however, about what would have happened to the subject had he or she been treated to some other level. For a subject who is exposed to level  $A = a'$ , all potential outcome  $\{Y_a\}_{a \in A}$ , except  $Y_{a'}$ , are unobserved and counterfactual. However, subject-specific causal effects are in general unidentifiable.<sup>3</sup>

A more useful concept is the population causal effect, which measures the aggregated effect of the treatment over the study population. Because the potential outcome  $Y_x$  may vary across subjects, we may treat it as a random variable that follows a probability distribution  $\Pr(Y_a)$ . In general,  $\Pr(Y_a)$  can be interpreted as the population proportion of subjects with an outcome equal to  $y$  under the hypothetical

---

<sup>3</sup>A rare exception is when we are able to observe the same subject under several treatment levels subsequently without any crossover effects. Under these situations, subject-specific causal effects can be identified (Piantadosi 2005; Morgan and Winship 2007).

scenario where everybody receives treatment level  $a$ . The population causal effect of treatment level  $a$  and  $a'$  is defined as a contrast between the potential outcome distribution  $\Pr(Y_a)$  and  $\Pr(Y_{a'})$ , for example causal mean difference  $E(Y_a) - E(Y_{a'})$ . When the outcome  $Y$  is binary, it would be natural to consider the causal risk ratio  $\frac{\Pr(Y_a=1)}{\Pr(Y_{a'}=1)}$  or the causal odds ratio

$$\frac{[\Pr(Y_a = 1)/\Pr(Y_a = 0)]}{[\Pr(Y_{a'} = 1)/\Pr(Y_{a'} = 0)]}.$$

If  $\Pr(Y_x)$  does not depend on  $a$ , then the treatment has no population causal effect on the outcome, and the causal null hypothesis holds. The converse is not true. It is logically possible that the treatment has a causal effect for some subjects, but that these effects ‘cancel out’ in such a way that there is no aggregated effects over the population.

Using potential outcomes, the fundamental difference between association and causation can be expressed clearly. In the population, some subjects are treated and some subjects are not. We say that treatment and outcome are associated in the population if the outcome distribution differs between the treated and the untreated. To quantify the association, we may, for instance, use the mean difference  $E(Y|A = 1) - E(Y|A = 0)$  or the risk ratio  $\Pr(Y = 1|A = 1)/\Pr(Y = 1|A = 0)$ . Thus, when we assess the treatment-outcome association, we are by definition comparing two different groups of subjects: those who are actually treated against those who are actually untreated. In contrast, the population causal effect compares the potential outcomes for the same subjects (the whole population) under two hypothetical scenarios: everybody being treated versus everybody being untreated. This

fundamental difference is the reason that association is in general not equal to causation in his framework. When we compare different subjects, there is always a risk that the subjects are different in other aspects than in the received treatment levels (Ho et al. 2007). If they are, then we may observe different outcome distributions for the treated and the untreated, even if treatment has no causal effect on the outcome.

In addition to the consistency or ignorability assumption, another important assumption is stable unit treatment value assumption (SUTVA).

**Assumption 2.** *Stable Unit Treatment Value Assumption (SUTVA): the observation on one unit should be unaffected by the particular assignment of treatments to the other units' (Cox 1958; Rubin 1980).*

Consider the situation with  $N$  units indexed by  $i = 1, \dots, N$ ;  $T$  treatments indexed by  $a = 1, \dots, T$ ; and outcome variable,  $Y$ , whose possible values are represented by  $Y_{ia}(a = 1, \dots, T; i = 1, \dots, N)$ . SUTVA is simply the a priori assumption that the value of  $Y$  for unit  $i$  when exposed to treatment  $a$  will be the same no matter what mechanism is used to assign treatment  $w$  to unit  $i$  and not matter what treatments the other units receive, and this holds for all  $i = 1, \dots, N$  and all  $a = 1, \dots, T$ .

According to Rubin, SUTVA is violated when unrepresented versions of treatment exist, e.g.,  $Y_{ia}$  depends on which version of treatment  $a$  is received, or when there is interference between units, i.e.,  $Y_{ia}$  depends on whether  $i'$  received treatment  $a$  or  $a'$ , where  $i \neq i'$  and  $a \neq a'$ . In clinical settings, SUTVA is, for example, violated when the members in the same family,  $A$  and  $B$ , are treated, and one of them,  $B$ , is



solely responsible for cooking. Although the effect of the treatment on a unit should not be affected by the effect of the treatment on other units, it is possible that the treatment on  $B$  might bias the treatment effect on  $A$  by affecting  $B$ 's taste buds. In policy evaluations, SUTVA is violated when a treatment alters social or environmental conditions that, in turn, alter potential outcomes. Winship and Morgan (1999, 663) illustrated this idea by describing the impact of a large job training program on local labor markets:

Consider the case where a large job training program is offered in a metropolitan area with a competitive labor market. As the supply of graduates from the program increases, the wage that employers will be willing to pay graduates of the program will decrease. When such complex effects are present, the powerful simplicity of the counterfactual framework vanishes.

SUTVA is both an assumption that facilitates investigation or estimation of counterfactuals and a conceptual perspective that underscores the importance of analyzing differential treatment effects with appropriate estimation.

As it turns out, SUTVA basically imposes exclusion restrictions. Heckman interprets these restrictions as the following two circumstances: 1) SUTVA rules out social interactions and general equilibrium effects and 2) SUTVA rules out any effect of assignment mechanism on the potential outcomes (Heckman 2005; Guo and Fraser 2010).

A recommended solution to SUTVA violation is, if possible, to change the unit of analysis to a higher level, at which the unit interference does not occur. For

example, if the violation of SUTVA is suspected at the individual level, we could change the unit of analysis to the household level. The problem is that this strategy is not always feasible in many observational studies, e.g., in the country-level analysis (Hong and Raudenbush 2013).

## Identification of Causal Effects

To reiterate, the consistency condition is expressed as follows:

$$\{Y_a\}_{a \in A} \perp\!\!\!\perp A \tag{2.2}$$

When Equation (2.2) holds, subjects are said to be exchangeable across treatment levels. Under consistency Equations (2.1) and (2.2), the conditional probability of  $Y$ , among those who actually received treatment level  $x$ , is equal to the probability of  $Y$ , had everyone received level  $x$ :

$$\Pr(Y = y|A = a) = \Pr(Y_a = y|A = a) = \Pr(Y_a = y) \tag{2.3}$$

The first equality in Equation (2.3) follows from Equation (2.1) and the second equality from Equation (2.2). Thus, under consistency and exchangeability, any measures of association between  $A$  and  $Y$  equals the corresponding population causal effect of  $A$  on  $Y$ . For example, the associational mean difference  $E(Y|A = 1) - E(Y|A = 0)$  equals the casual mean difference  $E(Y_1) - E(Y_0)$  and the associational relative risk  $\Pr(Y = a|A = 1)/\Pr(Y = 1|A = 0)$  equals the causal risk ratio  $\Pr(Y_1 = 1)/\Pr(Y_0 = 1)$ . Because randomization produces exchangeability, it follows that population causal effects are identifiable in randomized experiments.

Exchangeability means that all potential outcomes in  $\{Y_a\}_{a \in A}$  are jointly independent of  $X$ . Although this is a sufficient criterion for identification of the population causal effects, it is slightly stronger than necessary. By inspecting Equation (2.3), we observe that  $\Pr(Y_a)$  is identified for all  $a$  if the potential outcomes in  $\{Y_a\}_{a \in A}$  are separately independent of  $A$ :

$$Y_a \perp\!\!\!\perp A, \quad \forall a \tag{2.4}$$

In the literature, the word ‘exchangeability’ is sometimes used for the relation in Equation (2.4).

### Observational Studies

When the treatment is not randomized, exchangeability does not necessarily hold, and an observed association cannot in general be interpreted as a causal effect. Violations of Equation (2.2) typically occur when the treatment and the outcome have common causes. An illustration, suppose that we wish to study the effect of a policy program ( $A$ ) on outcome ( $Y$ ) for countries with certain problems. Suppose that a country’s general status affects what treatment level the country is assigned to (countries in a critical condition may, for example, receive higher treatments than countries in a noncritical condition). Moreover, a country’s status clearly affects its future outcome. That a country’s status affects both treatment and outcome implies that  $A$  and  $\{Y_a\}_{a \in A}$  are associated, which violate Equation (2.2). When the treatment and the outcome have common causes, we say that treatment-outcome association suffers from confounding. The standard way to deal with confounding is

to adjust for, i.e., condition on, a selected set of potential confounders, for example, by stratification or regression modeling. The rationale for this approach is that after adjustment, it may be reasonable to consider the treatment as being randomized by ‘nature.’ Formally, the aim of confounding adjustment is to produce conditional exchangeability.

$$\{Y_a\}_{a \in A} \perp\!\!\!\perp A|X, \tag{2.5}$$

where  $X$  indicates a set of covariates. Under consistency Equation (2.1) and conditional exchangeability Equation (2.5),  $\Pr(Y = y|A = a, X) = \Pr(Y_a = y|X)$ . It follows that any measures of the conditional association between  $A$  and  $Y$ , given  $X$ , equals the corresponding conditional population causal effect. For instance,  $\Pr(Y = 1|A = 1, X)/\Pr(Y = 1|A = 0, X)$  equals  $\Pr(Y_1 = 1|X)/\Pr(Y_0 = 1|X)$ . The population, not  $X$ -specific, causal effect can be computed through standardization, i.e., by averaging over the marginal confounder distribution.

$$\Pr(Y_a = y) = E\{\Pr(Y = y|A = a, X)\}$$

### Graphical Criteria for Identification

Potential outcomes framework focuses attention on whether conditional ignorability (a causal assumption) holds for a given set of adjustment variables. The major advantage here is that if conditional ignorability does hold given  $\mathbf{X}$  then adjustment for  $\mathbf{X}$  is guaranteed to be sufficient to control confounding. The major problem is

that the potential outcomes framework provides little guidance as to what sets of background variables are likely to produce conditional ignorability. Conditional ignorability is a global assumption that is defined for potential outcomes and is not strictly testable. By replacing this single large assumption with a series of local assumptions the deterministic structural equations models, we can get additional means of assessing the adequacy of various adjustment strategies.

The rules of Pearl’s *do*-Calculus give rise to a simple graphical criterion called the *back-door criterion* that can be checked to see if a given set  $\mathbf{Z}$  is sufficient to control confounding bias. Directed acyclic graph (DAG) is useful for illustrating causal relations (Pearl 2009; Edwards 2000; Koller and Friedman 2009; Lauritzen 1996). 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). In Pearl’s terminology, if there is a directed path from  $X$  to  $Y$  in a DAG,  $X$  is an ancestor of  $Y$ , and  $Y$  is a descendent of  $X$ .

**Definition 1.** (Back-Door Criterion (Pearl 2000, 79)) Given a causal model  $M$  and associated causal graph  $G_M$ , A set of covariates  $\mathbf{X}$  satisfies the back-door criterion for a causal variable  $A$  and outcome  $Y$  if:

1. no element of  $\mathbf{X}$  is a descendant of  $A$ ; and
2.  $A$  is  $d$ -separated from  $Y$  by  $\mathbf{X}$  in the graph  $G_A$  formed by deleting all edges out of  $A$  from  $G_M$ .

If  $\mathbf{X}$  satisfies the back-door criterion then the potential outcome distribution can be calculated using the standard stratification adjustment:

$$\Pr(Y_a = y) = \sum_{\mathbf{x}} \Pr(y|a, \mathbf{x}) \Pr(\mathbf{x}),$$

where  $\mathbf{x}$  may be multivariate. Pearl refers to this as the back-door adjustment. Since if  $\mathbf{X}$  satisfies the back-door criterion the standard stratification adjustment is appropriate, it follows that matching or stratifying on  $\Pr(a|\mathbf{x})$  (the propensity score given a realized value  $\mathbf{x}$  of  $\mathbf{X}$ ), along with related adjustments that make use of conditional ignorability, will also be appropriate (Rosenbaum and Rubin, 1983, 1984). As we will see below, this is true regardless of whether all (or even any) of the variables that affect treatment assignment are in  $\mathbf{X}$  all that is required is that conditional ignorability hold given  $\mathbf{X}$ . Again, the major advantage of this graphical approach to the identification of causal effects is that it is framed in terms of a series of local assumptions about causal mechanisms. These local assumptions are often easier to consider, debate, and possibly reject as unbelievable than the single global assumption of conditional ignorability.

Covariate adjustment is a technique capable of handling confounding in situations where sufficiently many potential confounders are observable. However, in practical causal inference problems, it is often the case that a non-causal path between treatment  $A$  and outcome  $y$  exists that consists solely of unobserved variables.

The relevant counterfactual restrictions implied in this graph are:  $(Y_{x,a} = Y_x, (X_a \perp\!\!\!\perp A), (Y_{x,a} \perp\!\!\!\perp X_a))$ . These restrictions can be used to produce the following

derivation:

$$\begin{aligned}
\Pr(Y_a) &= \sum_x \Pr(Y_a, X_a = x) \\
&= \sum_x \Pr(Y_a, x, X_a = x) \\
&= \sum_x \Pr(Y_a, x) \Pr(X_a = x) \\
&= \sum_x \Pr(Y_x) \Pr(X_a = x) \\
&= \sum_x \sum_{a'} \Pr(Y|x, a') \Pr(a') \Pr(X = x|a) \tag{2.6}
\end{aligned}$$

Here the first equality is by definition, the second by consistency, and third and fourth are restrictions, and the last by above restrictions used to repeat the derivations. In other words, this derivation expresses the causal effect of interest  $\Pr(Y_a)$  as a product of effects  $\Pr(X_a)$  and  $\Pr(Y_x)$  and then identifies each effect in this product separately.

It is possible to provide a graphical criterion for identification.

**Definition 2.** (Front-door criterion (Pearl 2000, 83)) A set of variable said to satisfy the front-door criterion to an ordered pair of variables  $(A, Y)$  if: (1)  $X$  intercepts all directed paths from  $A$  and  $Y$ , (2) there is no unblocked back-door path from  $A$  to  $Z$ , and (3) all back-door paths from  $X$  to  $Y$  are blocked by  $A$ .

One difficulty with using the front-door criterion in practice is that a multitude of counterfactual assumptions must hold. In particular, there must exist

observable variables that mediate every causal path from the effect variable  $A$  to the outcome variable and, moreover, those mediating variables must satisfy ignorability assumptions with respect to both effect and outcome variables. Nevertheless, one advantage of the front-door method of identification is that it gives an alternative way of handling if covariate adjustment or instrumental variable methods are unreasonable.

## **Interpreting Causal Inference from Decision-Theoretic Perspective**

I have so far discussed the definitions and conditions for causal inference from statistical viewpoint. The same framework of causal inference can be applied in natural and social science, because the framework is not basically agent-based. One of the unique features of human actions is that humans select and adjust in response to the environment and the changing situations, which I call a time-varying. Decision-theoretic framework can be useful in formalizing this aspect of human action, and can provide guidelines for subject-specific individualized treatment.

The causal statement that  $x$  causes  $y$  can easily be translated into decision-theoretic and rationalist terms, however. If  $x$  has a positive causal relationship with  $y$ , the rational decision-maker has to increase  $x$  in order to increase  $y$ .<sup>4</sup> To get the best decision rules, the decision-maker evaluates the effect of  $x$  by contrasting the outcomes for each scenario. The problem is, as in causal inference based on potential outcomes framework, only one outcome is realized for each subject, and other

---

<sup>4</sup>Elster might call it normative in this sense (Elster 1986).



outcomes are potential outcomes.<sup>5</sup> To generalize the decision-theoretic intuition, we need to begin with the simplest scenario: single actor decision making in a single stage.

### **Subject-specific Treatment and Decision Making**

Subject-specific treatment can be viewed as realization of certain decision rules; these rules dictate what to do in a given state of the subject. Thus, decision-theoretic notion, such as *utility*, can readily be adopted (von Neumann and Morgenstern 1980; Luce and Raiffa 1957).

### **Decision in Single-stage**

For simplicity, first consider a single-stage decision problem, where the decision-maker has to decide on the optimal treatment for an individual subject. Suppose the decision-maker observes a certain characteristic of the subject,  $o$ , and has to decide whether to prescribe treatment  $a$  or treatment  $a'$ , based on  $o$ . In this example, a decision rule could be: “give treatment  $a$  to the subject if his or her individual characteristic  $o$  is higher than a prescribed threshold, and treatment  $a'$  otherwise.” In other words, a decision rule is a mapping from currently available information (*state*) into the space of possible decisions.

Any decision can be evaluated in terms of its utility and the state in which the decision is made. Now let  $o$  denote the state,  $a$  denote a possible decision

---

<sup>5</sup>This is equivalent to the payoffs for the off-the-equilibrium path(s) in game theoretic terms (Bates et al. 1998).

(treatment), and  $U(o, a)$  denote the utility of taking the decision  $a$  while in the state  $o$ . The current statistical problem can be formulated in terms of the opportunity *loss* or *regret* associated with pair  $(o, a)$  by defining a loss function

$$L(o, a) = \sup_a U(o, a) - U(o, a'),$$

when the supremum is taken over all possible decisions for fixed  $o$ . The loss function is the difference between the utility of the optimal decision for state  $o$ , and the utility of the current decision  $a$  under the state. Clearly the goal is to find the decision that minimizes the loss function at the given state  $o$ ; this is subject-specific decision making since the optimal decision depends on the state. Equivalently, the problem can be formulated directly in terms of the utility without defining the loss function. In that case the goal would be to choose a decision so as to maximize the utility for the given state  $o$ . The utility function can be specified in various ways, depending on the specific problem. One of the most common ways would be to set to  $U(o, a) = E_a(Y|o)$ , i.e., the conditional expectation of the primary outcome  $Y$  given the state, where the expectation is computed according to a probability distribution indexed by the decision  $a$ . Alternatively, one can define  $U(o, a) = E(Y(a)|o)$ , where  $Y(a)$  is the potential outcome of the decision  $a$ .<sup>6</sup>

---

<sup>6</sup>Manski uses similar decision-theoretic framework for evaluation of social welfare programs (Manski 2007). In this framework, a welfare contrast is the difference between the utilities corresponding to two decisions, say  $a$  and  $a'$ , under the same state  $o$ , i.e.,

$$g(o, a, a') = U(o, a) - U(o, a')$$

Note that in the case where  $a$  is equal to the optimal decision, defined as the argument of the supremum of  $U(o, a)$ , the welfare contrast coincides with the loss or regret associated with  $a'$ .

Of course, this decision-theoretic formulation does not address the questions on the mechanisms by which the treatment works on the outcome. It does directly address the question of what is the effect of the causal action of treatment (or policy) on the outcome variable  $Y$ . It further addresses the important question of how this compares with the effect of the alternative action of not taking the treatment.

The quantity needed to solve the decision problem is

$$\text{ACE} := E_t(Y) - E_c(Y) \tag{2.7}$$

This is the decision theoretic explication of the concept of average causal effect (ACE) at single stage.

### **Decisions in Multi-stage**

Decision making problems often involve complex choices with multiple stages, where decisions made in one stage affect those to be made at another. In the context of multi-stage decisions, a dynamic treatment regime, also known as adaptive treatment, is a sequence of decision rules, one per stage of intervention, for adopting a treatment plan to the time-varying stage of an individual subject. Each decision rule takes a subject's individual characteristics and treatment history observed up to that stage as inputs, and outputs a recommended treatment at that stage. Recommendations can include treatment type, dosage, and timing. The reason for considering a dynamic treatment regime as a whole instead of its individual stage-specific com-

ponents is that the long-term effect of the current treatment strategy may depend on the performance of a future treatment plan.

In the current literature, a dynamic treatment regime is usually said to be optimal if it optimizes the mean long-term outcome, e.g., outcome observed at the end of the final stage of intervention.<sup>7</sup> The main goals in the area of multi-stage decision are 1) to compare two or more pre-conceived dynamic treatment regimes in terms of their utility, and 2) to identify the optimal dynamic treatment regimes, i.e., to identify the sequence of treatments that result in the most favorable outcome possible (highest utility).

Thus, any attempt to achieve these goals essentially requires knowing or estimating the utility functions (or some variations). Key notions from single stage decision problems can be extended to multi-stage decision without problems.

### **Methodological Implications for Longitudinal Studies**

For illustration, suppose that subjects are treated over two stages and can receive at each stage either treatment  $a$  or  $a'$ . If an individual was randomized to receive treatment  $a$  first and then  $a'$ , this individual will have a single observed outcome  $Y$  which corresponds to the potential outcome  $y$  under regime  $(a, a')$ , which we denote by  $Y(a, a')$ , and three unobservable potential outcomes:  $Y(a, a)$ ,  $Y(a', a)$ ,  $Y(a', a')$ , corresponding to outcomes under each of the other three possible regimes. As is clear even in this simple example, the number of potential outcomes and causal ef-

---

<sup>7</sup>However, at least in principle, other utility functions like the median or other quantities, or some other feature of the outcome distribution) can be employed as optimization criteria.

fects represented by contrasts between the potential outcomes can be very large, even for the moderate number of stages (Blackwell 2012). The optimal dynamic regime may be estimated while limiting the number of models specified to only a subset of all possible contrasts.

### **Time-varying Confounding and Mediation**

Longitudinal data present different challenges from cross-sectional data: 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  $Y$ . 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 variable  $X$ . It may be the case, then, that a variable is a confounder relative to one set of variable  $X$ , 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 to entirely to changes in  $O$ . In cross-sectional data, eliminating the bias due to a confounding effect is typically achieved by adjusting for the variable in a regression model.

Confounding in its simplest form can be visualized in a DAG if there is an arrow from  $O$  to  $A$ , and another from  $O$  to  $Y$ . 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-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. The outcome, measured at  $t_3$ , is denoted  $Y$ . Suppose there is an additional (unmeasured) variable,  $U$ , which is a cause of both  $O_2$  and  $Y$ .

Let me 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. Turn attention now to the effect of  $A_2$  on  $Y$ ;  $O_2$  confounds the 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, as  $U$  does in Figure 2.1. If one were to adjust for  $O_2$  in regression model, it would open what is called a back-door path in Pearl's terminology from  $Y$  to  $A_2$  via the path  $b-a-c-g$ . This is known as collider-specification bias, selection bias, Berksonian bias, Berkson's paradox, or, in some contexts, the null paradox.<sup>8</sup> Collider-stratification bias can also occur when conditioning on or stratifying by variables that are caused by both the

---

<sup>8</sup>This phenomenon was first described in the context of a retrospective study examining a risk factor for a disease in a sample from a hospital in-patient population in Berkson (1946). If a control group is also ascertained from the in-patient population, a difference in hospital admission rates for the control sample and case sample can result in a spurious negative association between the disease and the risk factor. For example, a hospital patient without diabetes is more likely to have cholecystitis, since they must have had some non-diabetes reason to enter the hospital in the first place.

Figure 2.1: A Two-stage DAG illustrating Time-varying Confounding and Mediating



treatment and the outcome.<sup>9</sup>

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 analysis to be explicit in modeling assumptions, particularly as the absence of an arrow between two variables (nodes) in 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.

### Assumptions for Causal Inference in Sequential Setting

As in static single-stage decision settings, a fundamental requirement for the potential outcomes framework is the *axiom of consistency*, which states that the

---

<sup>9</sup>Some argue for the need to distinguish selection bias and confounding. Selection bias refers to the bias caused by conditioning on post-treatment variables, while the confounding the bias caused by pre-treatment variables.

potential outcome under the treatment and the observed outcome agree. In other words, 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 same, regardless of the manner in which treatments are assigned. This is often plausible in medical treatments where it is easy to conceive of how to manipulate the treatments given to the patients, but less obvious for treatments that are modifiable by a variety of means.

Before introducing the necessary assumptions for estimating dynamic treatment regimes, let me introduce the following notations. Let  $\bar{a}_K \equiv (a_1, \dots, a_K)$  denote a  $K$ -sequence of treatments. Let  $(d_1, \dots, d_K)$  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  $\bar{O} \equiv (O_1, \dots, O_j)$  denote the collection of covariates observed up to stage  $j$  and  $\bar{A}_{j-1} \equiv (A_1, \dots, A_{j-1})$  denote the collection of past treatment prior to stage  $j$ . We can combine the treatment and covariate history up to the  $j$ th stage into a single history vector,  $H_j \equiv (\bar{O}_j, \bar{A}_{j-1})$ . To estimate dynamic treatment regimes from either randomized or observational data, two assumptions are required:

**Assumption 3.** *Stable Unit Treatment Value assumption (SUTVA): A subject's outcome is not influenced by other subjects' treatment allocation.*

**Assumption 4.** *No Unmeasured confounders: For any regime  $\bar{a}_K$ ,*



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

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

SUVTA is rarely violated in clinical setting. For example, it may be violated in special cases in clinical trials 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 (support group) where the interpersonal dynamics between group members could influence outcomes. However, its possible violation is, as I discussed in the earlier section, problematic in social science, in which interaction among subjects and the effects through general equilibrium are quite common.

The second assumption always holds under either complete or sequential randomization, and is sometimes called the sequential randomization assumption, sequential ignorability, or exchangeability, which is closely linked to the concept of stability. The assumption may also be 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 covariates history, at each stage the treatment actually received,  $A_j$ , is independent of future states and outcome under any sequence of future treatment,  $\bar{a}_j$ . That is, conditional on the past history, treatment received at stage  $j$  is independent of future potential covariates and

outcome:

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

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 censored subjects 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. Let  $\Pr_j(a_j|H_j)$  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

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

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 and realistic 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. Visual and bootstrap based approaches to diagnosing positivity violations have been proposed for one-stage settings. 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.

## Discussion

I have provided an overview of causal inference and potential outcome framework, with specific focus on the conditions of identification in observational longitudinal studies. Although the existing popular framework provides a useful toolkit for analyzing the causal effect of a treatment in static setting, it has limitations in applying to chronic problems that need long-term treatment. The main difficulty lies in the time-varying confounding of treatments and covariates, which show the interaction between the decision-maker and the environment. The existing statis-

tical tools are not well equipped to address these problems, and the econometric models, which are better equipped in modeling human actions, e.g., Heckman selection model, sometimes need strong distributional assumption. Theory of sequential decision-making in decision theory provides a useful conceptual framework for modeling this situations, although it is less prepared for estimating the treatment effects quantitatively.

## Chapter 3

### Causal Effects of Transitional Justice

#### Definitions of Transitional Justice

Transitional justice<sup>1</sup> is the link between the two broad concepts of *transition* and *justice* (Kritz 1995*a,b,c*; Teitel 2000, 2003). Although these two terms are ‘essentially contested’ (Gallie 1956; Collier, Hidalgo and Maciuceanu 2006; Connolly 1993; Landman 2006), they have specific meaning in transitional justice context. Although transition may refer to regime change of various kinds, which include democratic transition, negative or adverse transition, state failure, state demise, and state creation, among others, transition in transitional justice context usually refers to the one in liberal democratic direction. This definition inevitably excludes the transitional justice measures adopted without liberal democratic transition by authoritarian regime. The implications of this omission contribute to the biased estimation of the effect of transitional justice, because transitional justice and democratization tend to go together. Justice in transitional justice context is, according to the UN report, an ‘idea of accountability and fairness in the protection and vindication of rights and the prevention and punishment of wrongs’ (Anan 2004).

---

<sup>1</sup>Some authors point out the term ‘transitional justice’ is conceptually inaccurate because transitional justice does not indicate a particular kind of justice only applicable to transitional period (Posner and Vermeule 2004).

International Center for Transitional Justice (ICTJ) provides a comprehensive definition of transitional justice: the set of judicial and non-judicial measures that have been implemented by different countries in order to redress the legacies of massive human rights abuses. The UN Report on transitional justice defines it as “full range of processes and mechanisms associated with a society’s attempt to come to terms with the legacy of large-scale abuses, in order to ensure accountability, serve justice and achieve reconciliation” (Anan 2004). Note that both of these definitions are too comprehensive to be analytically useful. Transitional justice according to these definitions could range from the commonly used measures like international and domestic criminal prosecution, truth commissions (TRCs),<sup>2</sup> reparations, and compensations, to the informal measures such as official or unofficial apology to the symbolic measures for memorialization such as anniversary, monument or street naming, and to measures based on local tradition like *gacaca* court in Rwanda (Shaw, Waldorf and Hazan 2010; Stan and Nedelsky 2013a).

For the purpose of this dissertation, the following considerations are given in defining transitional justice. First, should we consider the measures adopted only in democratic transitions? Second, should we consider only the measures implemented by government or should we include the measures implemented by any governmental and non-governmental agents? Third, should we consider only the measures that target the practices of the previous regime or should we also consider the ongoing practices of human rights violations after democratic transition?

---

<sup>2</sup>Hereafter, I use truth commissions and TRCs (Truth and Reconciliations Commissions) interchangeably.

I take a minimalist definition of transitional justice: the measures adopted by the new government to deal with the past human rights abuses after democratic transition constitute transitional justice. By taking this definition, we exclude the following measures: 1) the measures adopted by various non-governmental organizations, 2) the measures adopted in countries that have not gone through democratic transition, and 3) the measures that are targeted at other issues than human rights abuses, e.g., corruption.

## **Varieties of Transitional Justice Mechanisms**

The conceptual continuum of transitional justice measures ranges from the Kantian deontological (maximalist) position that argues that prosecution should be pursued whenever possible to the utilitarian minimalist position that argues for formal or virtual immunity (Elster 1998). The early interests on transitional justice were mainly concerned with the determinants and the merits of punishing the perpetrators, i.e., criminal investigations and trials (O'Donnell and Schmitter 1986; Kritz 1996; Pion-Berlin 1994; Minow 1998; Snyder and Vinjamura 2003/4; Sikkink and Walling 2007; Thoms, Ron and Paris 2008).

The early scholars of democratization argued that trials for past human rights violations are politically untenable and likely to undermine new democracies and that if transitional justice is ever implemented, it should be quick and immediate after the transition (Huntington 1991; O'Donnell and Schmitter 1986; Nino 1996; Zalaquett

1992; Malamud-Goti 1996).<sup>3</sup> However, as more and more countries were democratized, a consensus quickly emerged that additional transitional justice measures to trials are necessary, and that *truth* and *justice* are mutually reinforcing and necessary, along with reparation and guarantee of non-repetition.<sup>4</sup> On the above conceptual spectrum, truth commission lies in the middle between the maximalist and the minimalist and is considered the second-best option to trials, and their use is advocated when trials may threaten stability of the new regime (Roht-Arriaza and Mariezcurrena 2006). A cynical position takes the view that truth commission is a “popular way for newly minted leaders to show their bona fides and curry favor with the international community” (Tepperman 2002, 128) to achieve other ends like enticing the foreign aids or entering international organizations like EU. According to trial advocates, however, truth commissions are not necessary when trials are available, and may even undermine justice unless they are used to build a case for future trials. Worse, some argue, truth commissions may provide perpetrators with a smoke screen for continued abuses (Snyder and Vinjamura 2003/4). Critics of truth commissions

---

<sup>3</sup>O’Donnell and Schmitter also suggested the difficulty of holding human rights trials in nascent democracies, although they admitted the possibility of human rights trials particularly where egregious human rights violations occurred. They conclude that “if civilian politicians use courage and skill, it *may not necessarily be suicidal* for a nascent democracy to confront the most reprehensible facts of its recent past” (O’Donnell and Schmitter 1986, 32, emphasis added). Notable here is that they rely on the politicians’ personal courage and skill, instead of *institutionalized mechanisms*, for successful implementation of transitional justice. Many human rights practitioners who participated in transitional trials have expressed similarly pessimistic views (Nino 1996; Zalaquett 1992; Malamud-Goti 1996).

<sup>4</sup>Despite the consensus that a holistic and multi-faceted approach is necessary, some tensions still exist among various mechanisms and their respective methods and aims. The sharpest chasm is about whether, under what conditions, amnesty can be granted for the wrongdoings of the past. The UN position, reflecting the widely held view, is that there should be no amnesty for genocide, crimes against humanity, war crimes and serious human rights violations.



fear that they may be dangerous because a commissioner's attempts at establishing a true record of past abuses may generate resentment among victims and perpetrators alike. Establishing painful "truths" in divided societies could provoke further tensions, inflaming volatile situations and providing new grievances to be exploited by cynical elites.


There is no fixed Procrustean measure of transitional justice that fits for all countries and circumstances without modification (Shaw and Waldorf 2010). However, the choice set that a new regime face can roughly be illustrated by Figure 3.1. The first choice for the new regime is whether to adopt transitional justice or not, and the common option is to hold a series of human rights trials. But trial is not the only option for the new government. The incoming governments can employ a range of options, sometimes in addition to or sometimes instead of trials, which include: (1) other sanctions such as lustration laws, bans, and purges, (2) investigations, such as truth commissions and independent inquiries, (3) reparation and rehabilitation, including government programs, (4) institutional reform, especially the establishment of human rights oversights and the introduction, restoration or amendment of constitution, and (5) immunity, via amnesties and pardons.

Punitive measures usually take the form of a criminal prosecution of the members of the former regime in international and domestic courts.<sup>5</sup> Criminal charges

---

<sup>5</sup>Although domestic trials have been widely used, countries increasingly adopt other forms of criminal prosecutions such as (1) prosecution in foreign, regional, or international courts, (2) *ad hoc* international criminal tribunals such as International Criminal Tribunal for the former Yugoslavia, (3) the International Criminal Court, and (4) hybrid courts. For detailed explanation of hybrid courts, which are emerging innovations in East Timor, Sierra Leone, and Kosovo (see Kritz 2004).

Figure 3.1: Choice for New Government on Transitional Justice



are typically filed against the individuals who were in charge of a military or police unit, an administrative branch, or the regime. Trials aim to restore justice by seeking individual accountability for their actions during public office.

The effectiveness of international (and domestic) tribunals typically depends on two factors: external factors related to commitments of and cooperation by states, and internal factors related to the functioning of the judicial institutions in the state. The ability of tribunals, domestic or international, to dispense justice largely depends on the balance of power, i.e., those who hold political and military power. The internal parameters concern the functioning of the actual judicial mechanism *per se*, thus respect for the due process of law, the security of lawyers and witnesses, the prosecutor's penal strategy, the determination of proof, the sentence and so on.

Transitional governments also employ other measures that are specifically intended to prohibit select or broad classes of individuals from participating in political affairs. One is *lustration*, a process of 'purification' that excludes various types of officials, functionaries and elites based on their actual or presumed complicity in past abuses. The measure originates from Eastern Europe, where a number of such laws

have been implemented, but have been increasingly adopted in other parts of the world to a varying degree. Extensive purges of the civil service, military, police and other key segments of government and society, as well as formal bans on discredited political parties, are used to similar effect.

Truth commission is the most widely used informational measure (Hayner 2001, 2010; Chapman and Ball 2001; Freeman 2006; Wiebelhaus-Brahm 2010).<sup>6</sup> Truth Commissions are institutional bodies set up to investigate the past history of human rights violations. Truth commissions usually (1) focus on the past, (2) investigate a pattern of abuses over a period of time, (3) are a temporary body, and (4) are officially sanctioned and employed by the state (Hayner 2001, 14). A truth commission issues an official report, which is an exemplar of a re-interpretative activity. Truth commissions, however, often do not confine its goals to re-interpretation of the past but also to proclamation of other goals in preambles of the reports: (1) to discover, clarify, and formally acknowledge past abuses, (2) to restore the dignity and facilitate the right to know of victims, (3) to contribute to justice and accountability, (4) to outline governmental relationship and recommend reforms, (5) to promote reconciliation and reduce conflict, and (6) to establish the legitimacy of the new regime (Hayner 2001). Truth commissions are temporary bodies, mostly with a mandate of six months to two years, and their temporal mandate is set in advance, often with the possibility of extension. The functions of truth commissions are 1) production of

---

<sup>6</sup>In cases where the transitional governments have been unwilling, unable or slow to undertake any formal measures, non-state actors have on occasion conducted their own independent inquiries. The primary feature that distinguishes these from truth commissions is the absence of official sanction and support.

the “truth”, 2) presentation of the “truth”, and 3) proclaiming of the fundamental values of the society.

Truth commissions and independent investigations usually provide bases for official reparations programs, though they are hardly prerequisites. The actual forms of reparations vary - monetary payments, restitution of property, job and pension reinstatement, special long-term benefits, etc. In part, these differences reflect past circumstances of violence and deprivation. Another consideration is the current needs of victims. Also, transitional governments exercise ultimate discretion over the choice of compensation and certain types may be more feasible given available resources. Land reform and property restitution, however, can present unique problems, because they may involve removing the current inhabitants and owners and the counterfactual calculation, depending on the duration of the previous regime (Waldron 1992; de Greiff 2006; Elster 2006; Gibson 2009).

Efforts to hold perpetrators accountable are often foreclosed by expansive amnesties and pardons. The typical mechanism is an act by the new legislature or a decree by the new president. In many countries like Argentina, Brazil, Chile, Nicaragua, and South Africa, however, amnesty laws actually were carried over from the previous regimes. These legacy statutes may even be retained by the new governments and upheld by the courts, as has been the case in Argentina and Chile, generally under the threat of violence from those who might otherwise be subject to prosecution. Likewise, pressure from the military in Argentina, Panama, and Uruguay led to amnesties, pardons and other formal constraints - introduced after regime change - that either cut short or overturned criminal prosecutions.

## Spread of Transitional Justice?

Table 3.1 shows the number of countries that adopted transitional justice measures ranging from human rights trials, truth commissions, amnesties, financial reparations, and lustrations from 1970 to 2008.<sup>7</sup> Not surprisingly, the most commonly adopted measure was amnesty; a hundred countries adopted it. Eighty seven adopted some form of human rights trials and fifty truth commissions. Thirty countries adopted reparation measures, and twenty seven lustrations.<sup>8</sup> Note that these include all countries that existed during the period, not just transitional countries. Some countries adopt transitional justice measures to deal with the historical injustices in the past, although they did not make democratic transitions recently. The examples include the transitional justice measure on historical injustices in U.S., Norway, Australia, New Zealand, and Northern Ireland, among many others (Barkan 2000; Stan and Nedelsky 2013*b*).

Figures 3.2 and 3.3 show the increasing trend toward the adoption of transitional justice mechanisms of various sorts since 1970.<sup>9</sup> Figure 3.2 indicates that the use of trials and truth commissions increased from around 2000, but the amnesty

---

<sup>7</sup>The data on transitional justice measures are based on my own coding of Amnesty International Annual Report, the country report of US Department of State, and Stan and Nedelsky (2013*b,c*). The resulting dataset was compared with the existing data, which include (Backer 2009; Olsen, Payne and Reiter 2010; Kim and Sikkink 2010).

<sup>8</sup>The reason that the countries that adopted reparations are small are mainly due to the fact that the datasets are based on the country description in the US state department and Amnesty International, among others. Usually reparation programs for the victims are not widely reported outside the given country.

<sup>9</sup>Sikkink and Walling (2007) call this phenomenon ‘justice cascade.’ However, it is not clear whether it is due to the sudden increase of the number of democratized countries or whether it really indicates justice cascade. It is an open question whether the relationship between the increase of the number of democratized countries and the justice cascade is linearly constant.

Table 3.1: The Number of Countries that Adopted Each Transitional Justice Measure

TJ Measures	Number of Countries that Adopted
Trials	87
Truth Commissions	50
Amnesties	100
Reparations	30
Lustrations	27

relatively decreased during the period. Note that Figure 3.3 shows the *cumulative* number of countries-years adopting each transitional justice mechanism over time. Notable here is the surge of truth commissions after 2000.

Sikkink and Walling (2007) claim that human rights trials improved physical integrity in Latin America by comparing the physical integrity at the transition year and ten years after transition for the countries that adopted human rights trials and those that did not in those ten years. Among the many problems that this simple comparison has, including no proper control, noteworthy is many countries adopted multiple transitional justice measures.

If we expand our time frame, we can see that many countries adopted a combination of multiple transitional justice measures, although they were not necessarily adopted at the same year. Table 3.2 shows that the number of countries that adopted each combination of transitional justice measures from 1970-2008. Common combinations include [Trial and Amnesty] (24), [None](16), [Trial, TRCs, Amnesties, Reparations](15), [Trial, TRC, Amnesty](15). Trials and amnesties are two of common elements in these combinations, except in [None].

Table 3.2: The Number of Countries in Each Combination of Transitional Justice Measures

TJ measures	Number of Countries
Trial, TRC, Amnesty, Reparation, Lustration	6
Trial, TRC, Amnesty, Reparation	15
Trial, TRC, Amnesty, Lustration	6
Trial, TRC, Amnesty	15
Trial, TRC, reparation, Lustration	1
Trial, TRC, Lustration	1
Trial, TRC	2
Trial, Amnesty, Reparation, Lustration	3
Trial, Amnesty, Lustration	8
Trial, Amnesty	24
Trial, Reparation	1
Trial	3
TRC, Amnesty, Reparation	1
TRC, Amnesty	3
Amnesty, Reparation	1
Amnesty, Lustration	2
Amnesty	14
None (Virtual Immunity)	16
Total	124

Figure 3.2: The Number of Countries that Adopted the Transitional Justice Year for the Particular Year (beginning and continuing), 1970-2008





Table 3.4 shows the number of years that each transitional countries adopted transitional justice measure in 1970-2008. Table 3.3 shows the descriptive statistics of the adopted years for each transitional justice measure. More accurately, the numbers in Table 3.4 indicate the years that each country adopted transitional justice measure. Argentina adopted trials in most years (24), and the average years is 4.645. Taiwan adopted most years of truth commissions (15), and the mean is 2 (median: 0). Albania adopted lustrations in most years (9), and Argentina reparations in



Figure 3.3: The Cumulative Number of Countries that Take the Transitional Justice, 1970-2008



most years (12), Uganda amnesties (16). The descriptive statistics show the heavy-tailed distribution for some measures such as reparation (2.927), TRCs (1.979), and lustration (4.792), although trials (1.267) also shows that the distribution might be skewed to the right (positively skewed).

Figure 3.4 shows the density plots of the key variables for the countries that adopted any transitional justice measures, except amnesties, i.e., trials, TRCs, lus-

Table 3.3: Years of Transitional Justice

	Trials	TRCs	Lustrations	Reparations	Amnesties
Minimum	0.000	0.000	0.000	0.000	0.000
1st Quartile	0.000	0.000	0.000	0.000	1.000
Median	3.000	0.000	0.000	0.000	3.000
Mean	4.645	2.000	0.419	1.008	4.194
3rd Quartile	7.000	3.000	0.000	0.000	6.250
Maximum	24.000	15.000	9.000	12.000	16.000

tration and reparations, and for those that did not.<sup>10</sup> Noticeable differences are found in polity score (democracy) and empowerment rights index.<sup>11</sup> The countries that adopted transitional justice measures tend to have higher scores in both democracy and empowerment rights, and, the density plots of the number of human rights non-governmental organizations (NGOs) and the independent judiciary index also show the countries that adopt transitional justice have slightly higher scores in both. Note that all of these are related to some aspects of democracy, and it is highly probable that they are collinear.

Table 3.4: The Number of Years that Each Country Adopted Transitional Justice Measures

Country	Trial	TRC	Lustration	Reparation	Amnesty
Afghanistan	6	0	0	0	11
Albania	5	0	9	0	5
Algeria	7	3	0	1	11
Angola	0	0	1	0	11
Argentina	24	2	2	12	10
Armenia	0	0	0	0	1
Azerbaijan	0	0	0	0	5
Bangladesh	2	0	0	0	11
Belarus	0	0	0	0	0
Benin	8	0	0	0	6

*continued on next page*

<sup>10</sup>TJ in the figure stands for transitional justice.

<sup>11</sup>An additive index summarizing government respect for electoral self-determination, domestic movement, foreign movement, religion, speech, assembly and association, and workers' rights.

Table 3.4, cont.

Country	Trial	TRC	Lustration	Reparation	Amnesty
Bolivia	13	3	0	1	3
Bosnia and Herzegovina	13	8	3	6	2
Brazil	0	0	0	2	1
Bulgaria	5	1	3	0	2
Burkina Faso	0	3	0	0	3
Burundi	1	5	0	0	6
Cambodia	7	0	1	0	3
Cameroon	0	0	0	0	2
Cape Verde	0	0	0	0	0
Central African Republic	3	6	0	0	7
Chad	8	3	0	1	12
Chile	19	5	0	7	4
Comoros	0	0	0	0	0
Congo	5	0	0	0	6
Croatia	13	0	0	0	2
Czech Republic	5	0	0	0	0
Czechoslovakia	3	0	2	1	1
Democratic Republic of the Congo	4	6	0	0	9
Djibouti	0	0	0	0	0
Dominican Republic	0	0	1	0	2
East Timor	9	7	0	0	1
Ecuador	5	2	0	10	3
El Salvador	5	2	0	11	9
Eritrea	1	0	0	0	1
Estonia	1	11	0	0	0
Ethiopia	12	1	0	0	10
Fiji	5	0	0	0	1
Gabon	2	0	0	0	4
Gambia	6	0	0	0	2
Georgia	0	0	0	0	3
German Democratic Republic	0	0	0	0	4
German Federal Republic	0	0	0	0	0
Germany	11	3	3	7	2
Ghana	4	3	0	1	7
Greece	7	0	1	0	5
Grenada	1	6	0	0	0
Guatemala	13	5	0	7	7
Guinea-Bissau	5	2	1	0	9
Guyana	0	0	0	0	0
Haiti	7	3	0	4	7
Honduras	16	1	0	2	7
Hungary	5	0	1	1	1
India	3	2	0	0	5
Indonesia	6	9	1	0	4
Iran	2	0	0	0	7
Iraq	9	0	7	1	12
Ivory Coast	2	0	0	0	6
Kazakhstan	0	0	0	0	0
Kenya	13	2	0	1	9
Kosovo	0	0	0	0	0
Kyrgyzstan	0	0	0	0	1
Laos	0	0	0	0	2
Latvia	2	0	0	0	0
Lebanon	5	11	0	0	2

*continued on next page*

Table 3.4, cont.

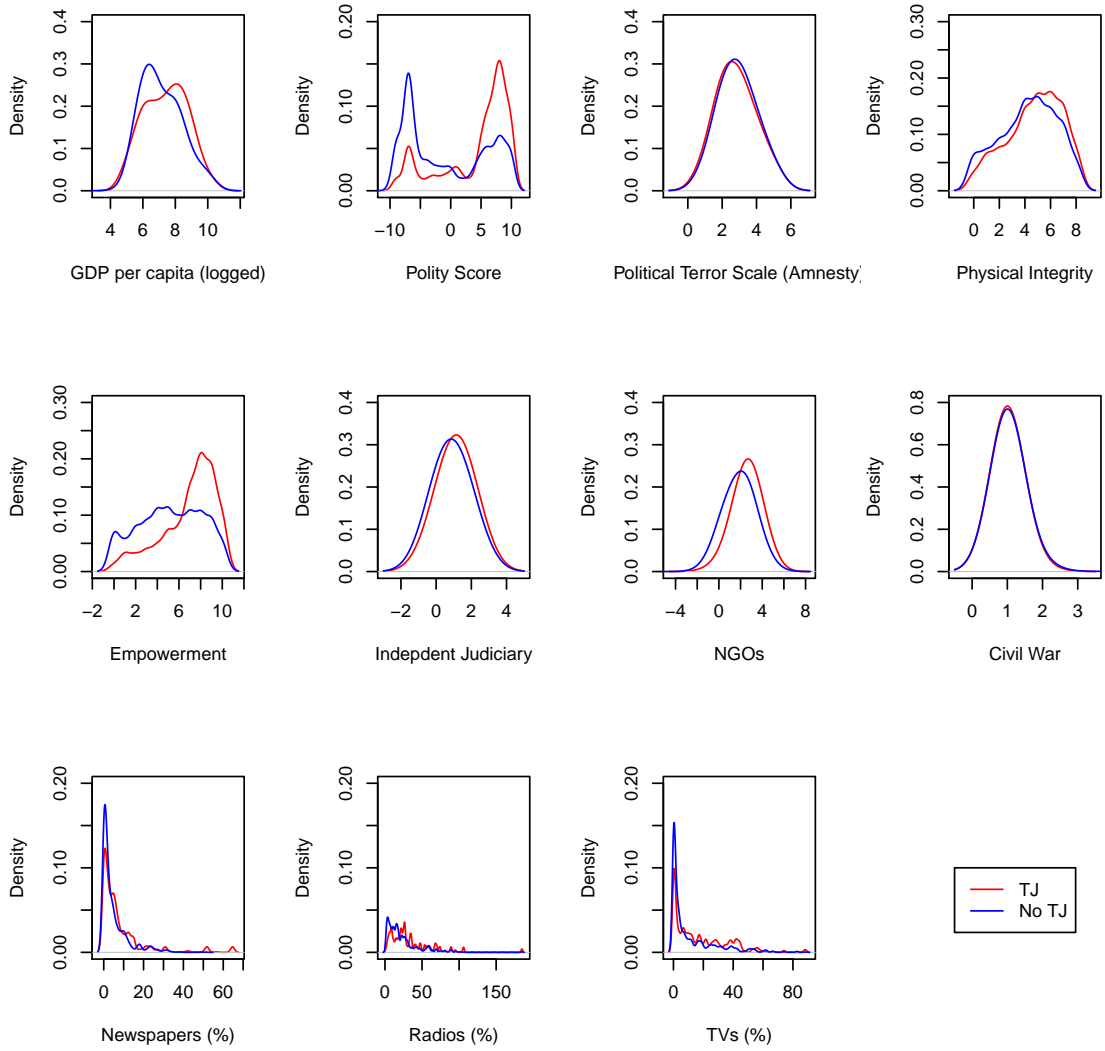
Country	Trial	TRC	Lustration	Reparation	Amnesty
Lesotho	1	2	0	0	2
Liberia	9	6	0	0	6
Lithuania	7	14	1	0	0
Macedonia	12	0	0	0	2
Madagascar	6	0	0	0	5
Malawi	3	0	0	1	1
Mali	11	0	0	0	3
Mauritania	3	0	0	0	6
Mexico	7	2	0	0	5
Moldova	0	0	0	0	3
Mongolia	0	0	0	0	0
Montenegro	0	0	0	0	0
Morocco	2	3	0	1	5
Mozambique	0	0	0	0	4
Namibia	2	0	0	0	1
Nepal	0	2	0	0	11
Nicaragua	15	0	0	9	8
Niger	3	0	0	0	9
Nigeria	5	4	2	0	7
Pakistan	2	4	0	0	5
Panama	14	3	1	0	3
Papua New Guinea	1	0	0	0	2
Paraguay	14	6	2	4	0
Peru	12	3	1	12	8
Philippines	4	2	0	2	13
Poland	20	3	1	1	4
Portugal	1	0	1	0	4
Romania	3	0	0	0	3
Russia	1	0	0	0	7
Rwanda	17	0	1	0	3
Sao Tome and Principe	0	0	0	0	1
Senegal	7	0	0	0	7
Seychelles	0	0	0	0	0
Sierra Leone	11	3	1	0	4
Slovakia	0	0	0	0	0
Slovenia	12	0	0	0	0
Solomon Islands	0	0	0	0	0
South Africa	17	12	0	4	10
South Korea	2	12	0	1	11
Spain	7	0	1	0	3
Sri Lanka	1	8	1	9	4
Sudan	2	0	1	0	14
Suriname	3	0	0	3	0
Taiwan	0	15	0	2	1
Tajikistan	0	0	0	0	3
Tanzania	0	0	0	0	2
Thailand	3	0	0	0	8
Togo	0	0	0	0	6
Tunisia	1	0	0	0	3
Turkey	1	0	2	0	5
Turkmenistan	0	0	0	0	0
Uganda	2	12	0	0	16
Ukraine	0	0	0	0	0
Uruguay	10	5	0	0	4

*continued on next page*

Table 3.4, cont.

Country	Trial	TRC	Lustration	Reparation	Amnesty
USSR	0	0	0	0	0
Uzbekistan	4	0	0	0	5
Yemen	0	0	0	0	3
Yugoslavia	7	2	0	0	5
Zambia	3	9	0	0	2
Zimbabwe	0	1	0	0	6

Figure 3.4: The Comparison of Density Plots of the Key Variables: The Countries that Adopted Transitional Justice and Those that Did Not



## Mechanisms of Transitional Justice Effects

Very few studies explicitly address the mechanisms of transitional justice effects or through what process transitional justice affect human rights. Two evident mechanisms can be found in the existing literature. One is concerned mainly with the perpetrators and the other victims, although the effects are not exclusive to them.

One view of the effects is top-down. The change of norms may trigger the change of incentive structure. Transitional justice measures can establish the norms of respects for human rights and socialize the *potential* perpetrators to respect them. This gradually socializes “ethnic cleansers” and warlords, for example, and eventually stabilizes the political system. Ideally, norms ‘corset’ states, progressively limits their room for maneuver, and control the behavior of the various players, especially potential spoilers, and imparts new values, gradually pervading the national institutions (Finnemore and Sikkink 1998; Hazan 2006).

While this mechanism largely emphasizes the effects of transitional justice as *norm-generator*, its logic is ultimately based on the cost-benefit calculation of repression (human rights violation) by the state officials. According to the rationalist explanation of repression, state officials choose repression because the benefits gained from repression exceed the costs (Poe and Tate 1994; Poe, Tate and Keith 1999). The expected benefits may include political gains from repressing political opponents, and the financial gains of expropriating their wealth and property. Trials lead to sanctions of various sorts (arrest, incarceration, loss of income and property) which may increase the perceived costs of repression for state officials. The main mechanism through which trials lead to human rights improvements is thus by increasing the

costs of repression for powerful state officials at the same time as the benefits of repression remain constant. Although this line of literature does not exclude social costs, the focus has been on material costs and benefits of repression. Hence the deterrence effect of transitional justice argument, which states that the increases in the probability or likelihood of punishment diminish repression.

Another view is bottom-up. Societies that have suffered from atrocities need to heal the victims as a pre-condition for reconciliation. The victims need to re-gain their dignity, recover the trust in the governmental institutions, and, be incorporated into the political system. This mechanism is concerned with emotional mechanisms based on catharsis and forgiveness. Its objective is national reconciliation, a process whereby former enemies manage to coexist without violence, which calls for a new societal pact to be drawn up which breaks the cycle of violence and vengeance (Minow 1998) - the fundamental tenets of transitional justice. It is based on the belief that victims are enabled to “regain” their identity by recounting their suffering publicly in the context of a criminal tribunal or truth commission (Hazan 2006).

## **Temporal Dimensions of Transitional Justice Effects**

The only empirical regularity closest thing a “law” in transitional justice is that ‘the intensity of the demand for retribution decreases both with time interval between the wrongdoings and the transition and with the interval between the transition and the trials’ (Elster 2004, 77). Although his intention is to emphasize the spatiotemporal variety of transitional justice measures rather than making a causal claim on the time interval and the public demand for transitional justice, we can be-



gin the discussion on the temporal dimension with this claim. Note that this claim is based on the individual emotional mechanism of anger that tends to fade over time. However, we could see many examples of the heightened demand for transitional justice, not immediately after the transition but after the new revelation of the past human rights abuses as in many East European countries. So, Elster's claim might be valid as far as the available information on the past abuses is constant over time, which is hardly the case in almost all transitional countries.

Different transitional justice mechanisms are expected to have different effects in given phases.<sup>12</sup> For example, we cannot expect the institutional reform can take effect immediately after implementation, while the effect of human rights trials could be instant. Taking temporal dimension is also important because the mechanisms of transitional justice not only reflect the balance of power but could also be instrumental in changing it. A consideration over time dimension reveals this dynamic reality. Defeated or weakened leaders often initially retain sometimes disruptive influences which may enable them to negotiate the form that the new internal order takes. However, their strength and ability to rally support tend to dwindle in the long run, mainly due to the effectiveness of transitional justice mechanisms.<sup>13</sup>

---

<sup>12</sup>The German example is particularly striking. Analyses show that until the 1960s most Germans saw the Allies tribunal in Nuremberg only as rendering the justice of victors. To their mind, the blanket-bombing of Dresden, Hamburg and Berlin by the US and UK air forces was the price already paid by German society for Nazi crimes. It was not until the 1970s that the Nuremberg Tribunal became an integral part of the German frame of reference and played a part in the younger generation questioning of their elders attitude during the war, a questioning reflected in the rapid rise of pacifism (Theissen 2008).

<sup>13</sup>This is what happened in Chile, where the truth commission contributed to the erosion of General Pinochet's popularity, promoting a new balance of power and opening the way, years later, for judicial proceedings which had formerly been impossible (Stan and Nedelsky 2013*b*).

Depending on the context, the effect of transitional justice is expected to be as follows in each phase (Hazan 2006).

**Repression/Armed Conflict phase** In this phase, the political and military leaders retain their partial or total control of power and make the work of international courts (the most widely used transitional justice mechanisms capable of intervening in this period) particularly difficult.

**Immediate post-transition term (first five years)** In this phase, warlords can (but do not necessarily) use their ability to cause disruption and can mobilize the media and networks loyal to them.

**Medium term (from five to twenty years)** In this phase, the society works out new points of reference. In the new political environment, those charged with offenses and the networks supporting them are weakened.<sup>14</sup>

**Long term** This phase is characterized by the rise of a new generation much more interested in overcoming old divisions.

Based on this timeline, Figure 3.5 shows an idealized progression of transitional justice over time. The primary effect in the first and the second phase is to weaken the collaborators of the old regime, or perpetrators, and the most effective transitional justice mechanisms are trials and sanctions, because those measures are

---

<sup>14</sup>The series of arrests in the former Yugoslavia, since more than 80 per cent of the accused — except for the two most famous fugitives — were taken into custody during this period.

targeting mostly the perpetrators. The indirect effect in these phases is demonstration effect of the commitment to the democratic principles, which is targeting the victims (and the non-collaborators of the regime) and the international community. The third and the fourth phases see institution-building and reconciliation. The direct effect is geared toward victims and the general population as opposed to the perpetrators. Note that all these processes are based on the assumption that the new government has a strong commitment to transitional justice, and that the effectiveness of the institutional-building and reconciliation are conditional on the effectiveness of trials and sanctions in the earlier period.

Figure 3.5: Timeline for the Ideal Progression of Transitional Justice


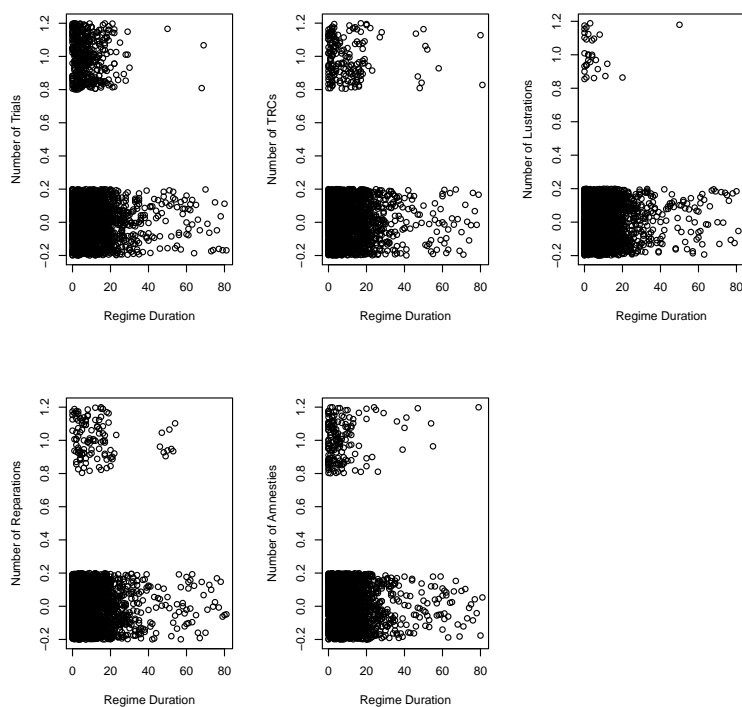


Figure 3.6 illustrates the (jittered) plots for regime duration and the adopted transitional justice measures. They clearly show that transitional justice measures like trials, reparation and lustration were adopted in the early period of the new regime, while truth commissions and amnesties were adopted even by the long lasting regimes. Note that these figures include all existing countries, transitional or non-transitional, and that, for example, TRCs include the commissions on historical injustices.

Figure 3.6: Adopted Transitional Justice Measures according to the Regime Duration



## Discussion

In this chapter, I discussed various aspects of transitional justice, adopted by the government. We could see several methodologically interesting points from the data. First, most of the countries have adopted more than one form of transitional justice, which indicates that treatments need to be modeled as multi-level, unlike the typical setting in causal inference literature where treatment is binary. Second, the distribution of adopted years of transitional justice measures shows that the distributions are skewed to the right, which means that some small number of countries

adopted transitional justice measures for protracted time, while many others did for a relatively short time. Third, the countries that adopted transitional justice measures like trials, truth commission, lustrations and reparations may have different characteristics from those that adopted none of these measures in terms of the level of democracy and human rights, among others.

A common problem in comparative politics is selection, which in this case means that the countries that adopt certain transitional justice measures already have different characteristics from those that chose other measures. As I discussed in Chapter 2, the relationship between the treatment and the outcome is associational, not causal, without properly adjusting the characteristics. In the case of transitional justice, it may be the case that the countries that adopt domestic human rights trials have more institutionalized judicial and democratic systems than those that do not choose human rights trials due to the lack of institutional resources. Furthermore, it is highly probable that the countries with human rights trials have experienced less human rights violations in pre-democratic transition period than those without human rights trials have because of the existence of the relatively developed judicial system and the democratic institutions. In the next chapter, I introduce the statistical tools to deal with this problem in various settings, and apply them to the problems.

## Chapter 4

# Using Propensity Score and Inverse Probability Weights

A commonly used tool in estimating for causal effect in observational studies is the propensity score. The propensity score is the probability of treatment conditional on observed covariates, and it can be used to balance covariates across treatment groups (Rosenbaum 2002, 2010; Guo and Fraser 2010). Propensity scores are typically estimated using parametric models such as the logistic regression model (for binary outcomes), and the individual propensity scores are typically compared using stratification or matching.

A big advantage of matching by propensity score is that it controls for many observed covariates simultaneously by matching subjects in one treatment group with subjects in another treatment group on the basis of individual propensity scores. The difference in average treatment effects between the two groups, typically the treated and the untreated, is calculated as the difference in outcomes between the matched groups. With stratification by propensity score, the average effect is calculated within each stratum, and the causal difference is estimated as the average of the within-stratum effects.

However, both matching and stratification have difficulty in some cases in con-

structuring comparison groups that may constrain their practical applicability. Matching may result in omitting a significant proportion of the population when comparison groups are being constructed because we can only focus on the overlapping groups (common support), thus limiting the generalizability from the results. Moreover, although stratification will produce treatment groups with similar probabilities for receiving treatment  $A$  and treatment  $B$ , the individuals in these strata may be indistinguishable. These problems are very common in political science, especially in comparative politics, where the group with one treatment, e.g., majoritarian electoral system, is different in many baseline characteristics from another group with a different electoral system, say proportional representation system, and the resulting overlapping observations after matching are very small.

Many scholars, notably Hirano and Imbens (2004), have recommended applying inverse propensity score estimators or inverse probability weighted (IPW) estimators to adjust for confounding, as a viable alternative to matching or stratification.<sup>1</sup> Semi-parametric inverse probability weighted estimators require fewer distributional assumptions about the underlying data than other methods, and they avoid the potential residual confounding that arises from stratification on a fixed number of strata. Inverse probability weighted estimators can also incorporate time-varying covariates and deal with the issue of censoring.

The purpose of this chapter is to introduce inverse probability treatment

---

<sup>1</sup>IPW estimators have a long history in statistics, beginning with the Horvitz-Thompson estimator, which was originally developed to deal with missing data in survey sampling (Horvitz and Thompson 1952; Lohr 2009).

weighting and inverse probability censoring weight, to extend it to the case of multiple treatments, and apply them in estimating the causal effects of transitional justice on human rights. In this chapter, I use inverse probability weight primarily to adjust for the sequence of multiple treatments, and the problem of time-varying confounding will be addressed in Chapter 5.

## Inverse Probability Weighted Estimators: Treatment and Censoring

Suppose we have a sample data from  $n$  subjects with treatments ( $A_i$ ), outcomes ( $O_i$ ), and individual covariates ( $X_i$ ), which are independent and identically distributed,  $1, \dots, n$ . The propensity score is usually not known and needs to be estimated from the observed covariates and treatment assignments. Denote the estimated propensity score as  $\hat{\pi}_a(\mathbf{X})$ , i.e., the estimates conditional probability of treatment given  $\mathbf{X}$ . The inverse probability weighted estimate of treatment specific effect,  $\mu_a$ , is given by the solution to the following estimating equation:

$$\sum_{i=1}^n \frac{I(A_i = a)(O_i - \mu_a)}{\pi_a(X_i, \hat{\gamma})} = 0, \quad a = \{1, 2\}, \quad (4.1)$$

where  $I(\cdot)$  is a treatment indicator function. A more widely known form of the average treatment effect is as follows: (Lunceford and Davidian 2004)

$$\text{ATE}_{\text{ipw}} = N^{-1} \sum_{i=1}^N \left( \frac{A_i Y_i}{\hat{\pi}_i(\mathbf{X})} \right) - N^{-1} \sum_{i=1}^N \left\{ \frac{(1 - A_i) Y_i}{1 - \hat{\pi}_i(\mathbf{X})} \right\} \quad (4.2)$$

In the above Equation 4.1, in order to estimate the causal effect of  $A_1$ , for



example, we need to include an observed outcome ( $O_i$ ) in the numerator if the individual received  $A_1$ .<sup>2</sup> The denominator of the estimating equation is the probability of receiving a given treatment, the propensity score. In cases of only two possible treatments, only a single propensity model needs to be fit, simply because an individual's probability of receiving  $A_2$  is one minus the probability of receiving  $A_1$ . What the equation ultimately shows is that observations with a high predicted probability for a given treatment receive a lower weight, compared with observations with a low predicted probability for a given treatment. Thus, an individual with a low predicted probability of receiving  $A_1$ , who actually received  $A_1$ , will represent a larger group of individuals who did not receive  $A_1$  (Robins, Rotnitzky and Zhao 1994; Curtis et al. 2007). The questions are when and why we need to use this estimators.

In regression models, average treatment effect is typically defined as

$$ATE = E(Y|A = 1, X) - E(Y|A = 0, X),$$

where the outer expectation is taken with respect to the distribution  $X$ . The empirical distribution of the conditioning set provides an easy estimate of  $F_x$  and simplifies the integration, so that the corresponding regression takes the form:

$$\widehat{ATE} = \frac{1}{n} \sum_{i=1}^n \{E(Y|A = 1, X_i) - E(Y|X = 0, X_i)\},$$

---

<sup>2</sup>Observed responses for individuals who received  $A_2$  are included only in the numerator of the estimate of the causal effect of  $A_2$ .

where  $E(Y|A = 1, X_i)$  is the estimated conditional expectation of the outcome given  $X_i$  within the treated group, and  $E(Y|A = 0, X_i)$  for the untreated group. Although these conditional expectation functions can be estimated using any consistent estimator such as OLS, GLM or GAM, question is that when the estimation is difficult over the full range of  $X$  due to high-dimensionality of  $X$ . When the observed values of  $X$  are dissimilar for the treated and the untreated, then either of  $E(Y|A = 1, X)$  and  $E(Y|A = 0, X)$  could be poorly estimated because the insufficient data points near either  $(A = 0, X)$  or  $(A = 1, X)$ . The estimation of over such non-overlapping ranges (non-common support) may underestimate the uncertainty in the estimator and/or result in finite sample bias. Inverse probability scores provide a reasonable solution in this situations (Glynn and Quinn 2010).

Inverse probability weighted estimators can also be used to deal with another thorny issue in observational studies, censoring, through inverse probability censoring weight (IPCW). In observational studies, the outcome, e.g., survival, is often observed after some period that may vary by individual subject. Because of this time lag and the limited follow-up in many studies, some outcome data will inevitably be right-censored because the study ends before the event, in this case death.

Denote the ascertainment time ( $T$ ), the potential censoring time ( $C$ ), and the treatment specific censoring distribution  $\{K_a(t)\}$ . As the true censoring distributions are typically not known, we also need to estimate them on the basis of observed data. Assuming no correlation between censoring and covariates, we can estimate the censoring distributions using Kaplan-Meier estimates stratified by treatment. Equation (4.1) can be expanded to the following:

$$\sum_{i=1}^n \frac{I(T_i < C_i)I(A_i = a)(O_i - \mu_a)}{\hat{K}_a(U_i)\pi_a(X_i, \hat{\gamma})} = 0, \quad a = \{1, 2\}, \quad (4.3)$$

where the treatment specific censoring distribution  $K_a(t)$  typically is estimated using Kaplan-Meier estimates of the censoring distribution. To the numerator we add an indicator variable to restrict the sample to uncensored individuals,  $I(T_i < C_i)$ , for whom the study ends before the censoring occurred. We expand the denominator to include a term reflecting the probability of staying in the sample (not being censored).

In general, however, censoring could depend on baseline and time-varying variables that, in turn, might be related to the subject's response to treatment. Therefore, a more conservative approach is to assume that the censoring process is conditional on covariate information.<sup>3</sup>

To summarize, the response of an uncensored individual to a given treatment is inversely weighted by the product of two probabilities: the probability of assignment to a given treatment and the probability of being uncensored. Hence, the observations that are more likely to be censored are weighted more heavily.

A problem of these weighted estimators is that only uncensored cases are included in the numerator. Suppose we are interested in estimating 5-year effects of a program, and an individual is lost to follow-up 4 years. The simple inverse probability weighted estimators require complete data and would exclude those 4 years of data. It seems reasonable that the information collected from partial ob-

---

<sup>3</sup>Under this assumption, we can estimate the censoring distribution using treatment-specific Cox proportional hazards models.

servations (four years) could be used to construct more efficient estimators. In this situation, partitioned estimators allow greater efficiency by incorporating data from these partial observations up to the point of censoring (Cole and Hernán 2008).

The general idea is to divide the follow-up period into non-overlapping partitions and estimate the average causal treatment effect within each partition. As a result of partitioning the follow-up interval, individuals considered to be censored for the simple weighted estimators may contribute their costs for one or more partitions. In general, partitioned estimators will have smaller asymptotic variance than the simple weighted estimators.

The censoring distribution within each partition can be modeled using a Kaplan-Meier estimator, but a more robust partitioned estimator can be constructed with Cox models. As with simple weighted estimators, the Cox version of the partitioned estimator is at least as efficient as the Kaplan-Meier version of the partitioned estimator based on the general theory of inverse probability-weighted estimators.

## **Propensity Scores and Inverse Probability Treatment Weights for Multiple Treatments**

In general, matching with multiple levels of treatment is an extension of propensity score matching under a binary condition. This method was originally proposed by Joffe and Rosenbaum (1999). When moving from a binary condition to multiple treatments, matching essentially requires the creation of matched pairs in such a way that high-level and low-level treatment groups have similar or balanced distributions of observed covariates. However, balancing covariates with propensity

scores under the condition of multiple doses raises three considerations. First, under a multiple treatment condition, the original definition of a propensity scores that the probability of the subject's receiving treatment given covariates,  $\Pr(Y_{a=1}|X)$ , is not applicable any more, because there are multiple treatments and each subject now has multiple propensity scores. If there are three levels of treatment, the subject has three propensity scores for each level of treatment. Joffe and Rosenbaum (1999) uses a single score balancing score and shows that such a score can be obtained in ordered logistic regression.<sup>4</sup>

Second, under the multiple treatment condition, one needs to re-define the distance between a treated case and an untreated case in optimization of matching. In this situation, the purpose is to identify the pairs that are similar in observed covariates, but very different in treatment levels. Here, the distance need to measure both the similarity in covariates and difference in treatment levels.

Third, under the multiple treatment condition, the matching algorithms adopted are different from those adopted in binary treatments. In the binary treatment setting, matching one group to another disjoint group is called the bipartite matching problem. But, matching under the condition of multiple treatment levels is matching within a single group and a non-bipartite matching.

I introduce two methods to deal with multiple treatment; one suggested by Joffe and Rosenbaum (1999) and the other by Imbens (2000). Joffe and Rosenbaum

---

<sup>4</sup>Using OLS could be problematic because OLS assumes constant error variance, but error variances might vary by the treatment levels, and possible heteroskedasticity (Guo and Fraser 2010, 164)

(1999) suggest the following procedure:

1. Develop a single scalar score based on an ordered logistic regression. The distribution of treatment levels  $Z_K$  for a sample of  $K$  subjects ( $k = 1, 2, \dots, K$ ), given observed covariates  $\mathbf{x}_k$ , is modeled as

$$\log \left( \frac{\Pr(Z_k \geq d)}{\Pr(Z_k < d)} \right) = \theta_d + \boldsymbol{\beta}' \mathbf{x}_k, \quad \text{for } d = 2, 3, 4, 5,$$

assuming there are five treatment levels to be modeled. This model compares the probability of a response greater than or equal to a given category ( $d = 2, \dots, 5$ ) to the probability of a response less than this category, and the model is composed of  $d - 1$  parallel linear equations. In this model,  $\theta_d$ , cutoff value, is used in calculating predicted probabilities of each of the five responses.

2. Calculate the distance between participants  $k$  and  $k'$ , where  $k \neq k'$ . Lu et al. (2001) provide the following equation to calculate the distance under the multiple doses condition:

$$\Delta(\mathbf{x}_k, \mathbf{x}_{k'}) = \frac{(\hat{\boldsymbol{\beta}}' \mathbf{x}_k - \hat{\boldsymbol{\beta}}' \mathbf{x}_{k'})^2 + \varepsilon}{(Z_k - Z_{k'})^2} \quad (4.4)$$

where  $\boldsymbol{\beta}' \mathbf{x}_k$  and  $\boldsymbol{\beta}' \mathbf{x}_{k'}$  are the estimated propensity scores, and  $Z_k$  and  $Z_{k'}$  are dose values ( $1, 2, \dots, d$ ) if there are  $d$  doses for  $k$  and  $k'$ , respectively.

3. Conduct non-bipartite pair matching using the distance as defined above. For a sample of  $K$  participants, each participant  $k$  has a distance from each of the

remaining participants in the sample on the estimated propensity scores. The researcher then conducts an optimal pair matching in such a way that the total distance associated with all matched pairs is minimized. Each of the resultant pairs then contains one high-dose participants and one low-dose participants, because  $\Delta(x_k, x_{k'}) = \infty$  if  $Z_k = Z_{k'}$ , which is forbidden by Equation (4.4).

4. Check covariate balance after matching. Having obtained matched pairs, the next step involves checking covariate balance between high- and low-dose participants to see how well the propensity score matching performed.
5. Evaluate the impact of treatment doses on the outcome.

The second method to deal with multiple treatments is suggested by Imbens (2000). He proposed to estimate multiple balancing scores by using a multinomial logit model, and then conduct an outcome analysis that employs the inverse of a specific propensity score as sampling weight. This method requires fewer assumptions and is easier to implement. Imbens's method could be used with several unordered treatments. This approach has two steps: 1) Estimate the generalized propensity score by using the multinomial logit, and 2) conduct outcome analysis by following propensity score weighting. In the first step, each subject receives the propensity scores for each treatment level. Hence, unlike binary treatment, the propensity scores are multiple and not scalar functions. The second step is to calculate the inverse probability treatment weights and use them as sampling weights in the analyses.

## **The Effect of Human Rights Trials on Human Rights**

Many countries simultaneously have to deal with many developmental problems like “poverty, inequalities, weak institutions, broken physical infrastructure, poor governance, high levels of insecurity, and low-levels of social capital, among many others” (Arenhövel 2008; de Greiff and Duthie 2009) — many of them are considered to have higher priority than dealing with the past for many countries, especially for those emerging from civil conflicts. As many political leaders and the practitioners of human rights claimed, “justice is not cheap” (IRIN 2006). The question of which transitional justice measure(s) are more effective has a practical implication for those countries, given the accompanying costs for transitional justice measures. As I discussed in Chapter 3, there is a significant difference between the maximalist and minimalist on transitional justice’s effects on human rights (Olsen, Payne and Reiter 2010). The previous studies have produced mixed results on the effects of human rights trials on human rights and repression (Kim and Sikkink 2010; Olsen, Payne and Reiter 2010).

The empirical analysis in this chapter begins with the observation that many countries adopted multiple transitional justice measures, either simultaneously or sequentially. Then, the better empirical question to ask is which combinations are most effective compared to others than whether, for example, human rights trials are more effective than TRC’s, because many countries adopt multiple measures over time.

Causal arguments are inevitably mechanism dependent and causal claims without mechanisms are not possible. A well-known mechanisms for the effect of



transitional justice measures on human rights are deterrence effects, expressed in terms of rationalist explanations. According to this explanation, state officials choose repression because the benefits gained from repression exceed the costs (Poe and Tate 1994; Poe, Tate and Keith 1999). The expected benefits may include political gains from repressing political opponents, and the financial gains of expropriating their wealth and property. Trials lead to sanctions of various sorts like arrest, incarceration, loss of income and property, which may increase the perceived costs of repression for state officials. The main mechanism through which trials lead to human rights improvements and less repression is thus by increasing the costs of repression for powerful state officials at the same time as the benefits of repression remain constant.

## **Data**

The data consist of the countries that made democratic transitions in 1970-2008. We are confronted with the problems in analyzing panel datasets: unbalanced panels and censoring. Unbalanced panels are due to the fact that countries in the data made transitions in different years. Some countries made democratic transitions in the mid-1970s, e.g., Spain, and the democratic regime has lasted more than thirty years, while some countries, like Bhutan or Guinea-Bissau, have less than five years of democratic regime, counting from the most recent democratic transition.

Some countries make multiple transitions; some of them are in the democratic direction, and some are in the non-democratic direction (authoritarian or state failure). Typical practice in the field is to identify the year that a country made demo-

cratic transition, whether it be from authoritarian regime, civil war or state creation, and take the year as the year zero of post-transition. This is not problematic for most countries with single democratic transition in the sample period like Argentina and South Africa, which have 1983 and 1994 as the years of democratic transition, respectively, and have no other transitions since then. Problematic is the countries that made multiple transitions. Typically, those have experienced major democratic transition(s), negative or adverse transition, or state failure. For example, take the cases of Guinea-Bissau and Haiti.

Figure 4.1: Single and Multiple Transitions

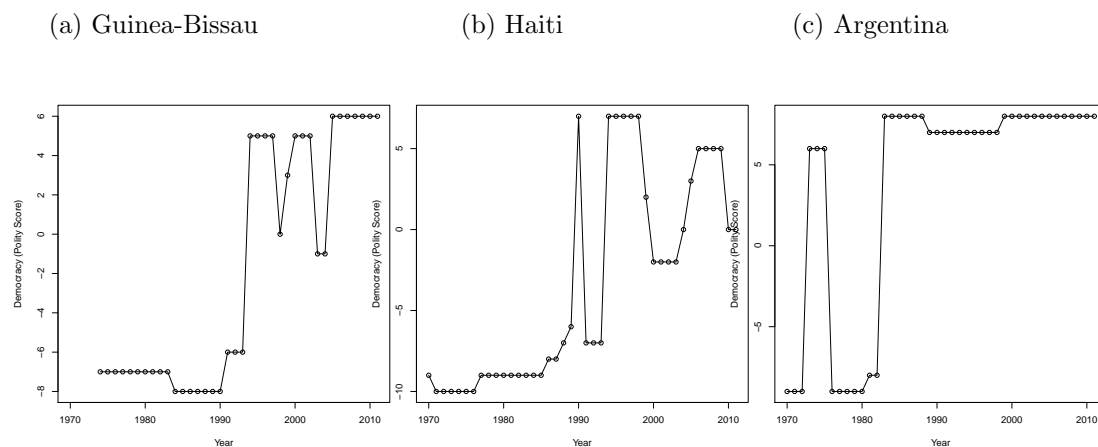


Figure 4.1 shows the polity scores of Guinea-Bissau, Haiti, and Argentina in 1970-2011. Unlike Argentina in Figure 4.1c, which has been democratic since 1983 without interruption, Guinea-Bissau and Haiti experienced multiple regime changes.<sup>5</sup> Negative numbers indicate authoritarian regimes, and positive numbers democratic

<sup>5</sup>The years of 2010 and 2011 for Haiti are coded neutral (0) in Polity IV due to state failure (-77).

ones. Based on Polity IV scores, we can identify six major transitions in Haiti during this period: three democratic (1986-1990, 1994, 2004-2006) and three adverse (1991, 1999 and 2000). We can identify three democratic spells, 1990, 1994-1999, and 2005-2009, and at least three authoritarian spells, 1970-1989, 1991-1993, and 2000-2004. The typical approach is to consider the first transition, 1986-1990, as the cutting point, and 1991 as the year zero of post-transition, and to consider the following years (1992-present) as if they were a single democratic regime, despite the many non-democratic intervening spells during the period.

Third problem is due to censoring. If we estimate the effects of human rights trials in five years or ten years, the democratic regimes that collapsed before the evaluation years, fifth or tenth year, are censored from the sample, hence missing from the data. This may contribute the overestimation of the transitional justice's effects.

As I defined in Chapter 3, transition is typically defined as a regime change in democratic direction, and includes transition to a democratic regime from an authoritarian one, transition from civil war, and state creation.<sup>6</sup> Table 4.1 shows the numbers of the regimes and the democratic regimes for the countries in the dataset since 1970. Guinea-Bissau experienced most transitions (12), and the mean number of regimes is 2.68 (median: 2). Note that more than a half of the countries in the dataset (57%) have just one or two regimes. One is typically for state-creation, e.g., post-communist countries, and two for a single democratic transition. It also

---

<sup>6</sup>`regtran` variable in the POLITY IV dataset was used to identify the country-years for transition.

clearly shows the heavy-tailed distribution, and raises problems in our analysis. If we restrict our analysis to democratic regimes, the regimes with short democratic lives are dropped from the sample. In other words, the countries with five years of democratic spell have smaller windows of opportunity to adopt transitional justice than those with ten years of democratic regime.<sup>7</sup>

Table 4.1: The Number of Regimes and Democratic Regimes

Country	Number of Regimes	Number of Democratic Regimes
Haiti	6	3
Dominican Republic	3	2
Mexico	4	2
Guatemala	4	3
Honduras	2	2
El Salvador	2	2
Nicaragua	4	1
Panama	2	1
Guyana	3	2
Suriname	5	2
Ecuador	3	3
Peru	4	3
Brazil	2	1
Bolivia	4	1
Paraguay	3	1
Chile	3	2
Argentina	4	2
Uruguay	3	1
Spain	2	1
Portugal	2	1
Germany	2	1
German Federal Republic	1	1
German Democratic Republic	1	0
Poland	2	1
Hungary	2	1
Czechoslovakia	2	1
Czech Republic	1	1
Slovakia	1	1
Albania	2	1
Montenegro	1	1
Macedonia	2	2
Croatia	2	1
Yugoslavia	4	2

*continued on next page*

<sup>7</sup>One of the possible ways to deal with this problem — uneven windows of opportunity to adopt transitional justice — is to use the years under democracy as an exposure or an offset variable, although this implicitly requires linearity assumption between exposure and treatments in the sense that two out of ten is equivalent to one out of five (Cameron and Trivedi 2009, 2013).

Table 4.1, cont.

Country	Number of Regimes	Number of Democratic Regimes
Kosovo	1	1
Slovenia	1	1
Greece	2	1
Bulgaria	2	1
Moldova	1	1
Romania	2	1
USSR	1	1
Russia	1	1
Estonia	1	1
Latvia	1	1
Lithuania	1	1
Ukraine	1	1
Belarus	2	1
Armenia	3	2
Georgia	1	1
Azerbaijan	2	1
Cape Verde	2	1
Guinea-Bissau	11	3
Gambia	2	1
Mali	2	1
Senegal	3	1
Benin	3	1
Mauritania	2	1
Niger	4	3
Ivory Coast	2	2
Burkina Faso	4	2
Liberia	3	2
Sierra Leone	3	3
Ghana	7	4
Togo	2	0
Cameroon	2	0
Nigeria	4	1
Gabon	3	1
Central African Republic	3	1
Chad	3	0
Congo	3	1
Democratic Republic of the Congo	2	1
Uganda	5	1
Kenya	3	1
Tanzania	2	0
Burundi	3	1
Rwanda	2	0
Djibouti	2	1
Ethiopia	2	1
Eritrea	1	0
Angola	2	0
Mozambique	2	1
Zambia	5	3
Zimbabwe	5	3
Malawi	2	1
South Africa	2	2
Namibia	1	1
Lesotho	3	2
Madagascar	6	3

*continued on next page*

Table 4.1, cont.

Country	Number of Regimes	Number of Democratic Regimes
Comoros	5	4
Morocco	1	0
Algeria	5	1
Tunisia	2	0
Sudan	4	1
Iran	4	1
Turkey	5	3
Iraq	2	1
Lebanon	2	2
Yemen	1	0
Afghanistan	3	0
Turkmenistan	1	0
Tajikistan	3	0
Kyrgyzstan	1	2
Uzbekistan	1	0
Kazakhstan	1	0
Mongolia	2	1
Taiwan	2	1
South Korea	3	2
India	1	1
Pakistan	6	3
Bangladesh	7	3
Sri Lanka	1	1
Nepal	5	2
Thailand	6	5
Cambodia	4	2
Laos	1	0
Philippines	3	2
Indonesia	2	1
East Timor	1	1
Papua New Guinea	1	1
Solomon Islands	2	2
Fiji	5	3

### Analysis 1: Moving Beyond Binary Treatment

The outcome variable in the model is physical integrity rights, a band of basic human rights measured at the country-level. Physical integrity rights typically include freedom from disappearance, extrajudicial killing, torture and political imprisonment. Political Terror Scale (PTS) is used as the measures for country-level physical integrity rights.<sup>8</sup> Its coding is based on the country description of hu-

<sup>8</sup><http://www.politicalterroryscale.org/>. For the discussion on the difference between PTS and another widely used measure, Cingranelli and Richards Human Rights Index (CIRI Index),

man rights conditions for the given year available in U.S. Department of State and Amnesty International, whose coding scheme is as follows:

1. Countries under a secure rule of law, people are not imprisoned for their view, and torture is rare or exceptional. Political murders are extremely rare (1).
2. There is a limited amount of imprisonment for nonviolent political activity. However, few persons are affected, torture and beatings are exceptional. Political murder is rare (2).
3. There is extensive political imprisonment, or a recent history of such imprisonment. Execution or other political murders and brutality may be common. Unlimited detention, with or without a trial, for political views is accepted (3).
4. Civil and political rights violations have expanded to large numbers of the population. Murders, disappearances, and torture are a common part of life. In spite of its generality, on this level terror affects those who interest themselves in politics or ideas (4).
5. Terror has expanded to the whole population. The leaders of these societies place no limits on the means or thoroughness with which they pursue personal or ideological goals (5).

---

see Cingranelli and Richards (1999); Wood and Gibney (2010); Cingranelli and Richards (2010). I chose PTS over CIRI for two reasons: 1) PTS covers more years than CIRI. CIRI is available from 1980, while PTS from 1976, and 2) PTS is less prone to yearly fluctuations than CIRI is, because CIRI emphasizes the practices, while PTS the conditions.

The reversed Amnesty scores in PTS are used as measures of physical integrity throughout the analysis, and higher scores indicate better physical integrity rights. The main explanatory variable of interest is transitional justice measures adopted by each democratic regime, which include human rights trials, truth commissions, and official amnesties.<sup>9</sup> The dataset on transitional justice measures is in the format of country-year absence/presence of prosecutions, truth commissions, pardons, etc (1970-2008), based on the country description of human rights practices.

As a starting point, I compare the mean scores of physical integrity scores at the transition year (year 0), the third, the fifth, and the tenth year after democratic transition. Note that the unit of analysis is the regime-year, not the country-year, and the same country may have multiple democratic regimes. Table 4.2 shows the mean physical integrity scores of the countries that adopted human rights trials, truth commissions, and amnesties. Note that if a country degenerates into non-democratic regime before reaching the milestone years, it is dropped from the computation.

First, Table 4.2 shows that the democratic regimes that adopted trials had higher scores than the ones that did not at baseline (transition years) and that the improvement for the regimes with trials are smaller than that for those without trials. Truth commissions tend to be adopted in regimes with worse human rights conditions. The stark difference can be seen in amnesties. The regimes with amnesties clearly have lower baseline values and smaller increases over the years.

---

<sup>9</sup>I exclude lustration and reparation measures here because relatively fewer countries with regional concentration adopted those measures.



Table 4.2: Physical Integrity Before and After Transitional Justice: 3, 5, and 10 years after Democratic Transition. The numbers in the parentheses indicate the number of democratic regimes used for computation of the mean scores.

		year 3		year 5		year 10	
		pre (136)	post (98)	pre (136)	post (82)	pre (136)	post (60)
Trials	no	3.049	3.300	3.054	3.500	3.123	3.474
	yes	3.236	3.354	3.210	3.320	3.127	3.342
TRCs	no	3.149	3.313	3.149	3.444	3.170	3.415
	yes	3.057	3.355	3.057	3.286	3.057	3.294
Amnesty	no	3.422	3.711	3.431	3.862	3.446	3.850
	yes	2.861	3.000	2.897	3.132	2.900	3.150

The before and after comparison, which is largely similar to Sikkink and Walling (2007), has a serious problem. As I mentioned in Chapter 3, many countries adopted multiple transitional justice measures. Given that we are interested in only three mechanisms of trials, truth commissions, and amnesties, eight combinations of treatments are possible in principle: (None), (Trials), (TRCs), (Amnesties), (Trials, TRCs), (Trials, Amnesties), (TRCs, Amnesties), (Trials, TRCs, Amnesties).<sup>10</sup>

Table 4.3 shows the mean physical integrity scores at year 0, year 3, year 5, and year 10 after democratic transition according to the adopted transitional justice combinations. For example, the first row in the first column indicate that the regimes that adopted none of the three transitional justice measures in the first three years of democratic transition have the mean scores of 3.389 at baseline, and 3.500 at year

---

<sup>10</sup>Olsen, Payne and Reiter (2010) use ordered categories composed of three levels: trials, TRCs, and amnesties. The country-years that adopted multiple transitional justice measures are categorized into the highest category. For example, a country that adopted TRCs and amnesties are in the same category as TRCs only (2), and a country which adopted trials and truth commission are in the same category with the country that adopted only trials (3).

Table 4.3: Physical Integrity Before and After Transitional Justice: Combined Measures. Numbers in the parentheses are standard errors.

	Year 3 (95)		Year 5 (76)		Year 10 (59)	
	Pre	Post	Pre	Post	Pre	Post
None	3.389(.257)	3.500(.294)	3.250(.479)	3.500(.866)	3.400(.510)	3.400(.510)
Trials	2.842(.233)	2.947(.259)	2.800(.243)	3.000(.169)	3.600(.245)	3.200(.200)
TRCs	4.667(.333)	4.000(.000)	4.500(.500)	4.000(.000)	4.500(.5)	4.000(1.000)
Amnesties	3.444(.294)	3.111(.351)	3.571(.369)	3.286(.286)	3.667(.333)	3.667(.422)
Trials, TRCs	2.286(.184)	2.571(.202)	2.667(.333)	2.667(.333)	2.667(.667)	2.000(.000)
Trials, Amnesties	3.125(.221)	3.063(.193)	3.000(.192)	2.950(.246)	3.045(.180)	3.000(.186)
TRCs, Amnesties	3.625(.420)	3.750(.313)	3.333(.494)	3.500(.342)	3.500(1.500)	4.000(1.000)
Trials, TRCs, Amnesties	2.667(.303)	3.333(.319)	2.875(.301)	3.313(.270)	3.143(.312)	3.571(.251)

3. The biggest increase can be seen in the regimes that adopted all three transitional justice measures.

## Analysis 2: Moving to the Inclusion of Covariates and Adjustment

The simple comparisons in the previous sections did not include any controls. In this section, I include structural variables in addition to transitional justice variables in addition to transitional justice variable.

- Lag dependent variable<sup>11</sup>

---

<sup>11</sup>The effect of including the lagged dependent variable in the right side of the equation is notoriously well-known (Achen 2001; Keele and Kelly 2006). Some argue for its inclusion as it typically mitigates to a very large extent any problems with autocorrelation in the data. Theoretically, the lagged dependent variable should be included if human rights performance in one year truly affects human rights performance next year. This could be justified if, for example, there is a reason to presume that a history of applying torture makes state officials accustomed or habituated to the application of torture. In such cases, even if torture were to become formally punished by the ruling political authorities, this might not affect a change in actual behavior by lower tier state officials or might affect a change only with a substantial delay. Again the inclusion of a lagged dependent variable, speaking that it typically absorbs an enormous amount of variation in the dependent variable, leaving little for the remaining independent variables to explain as well as sometimes leaving coefficients with the wrong sign. In line with the existing studies, I include a lagged dependent variable in the models to be estimated with caution. Note that this could lead to Nickell bias in the estimation, which for large  $N$  becomes smaller as  $T$  increases, however (Nickell 1981).


- **polity**: Polity Score ranging from -10 to 10 (Polity IV)
- **lngdp**: log GDP per capita (World Bank)
- **growthrate**: Economic Growth Rate (World Bank)
- **lnpop**: logged population (World Bank)
- **popgrowth**: population growth rate (World Bank)
- **NGOs**: total number of human rights NGOs involved in a country (internally) in a given year (Mosley and Uno 2007)
- **treaties**: number of the key human rights treaties (out of three) that the country ratified (0-3)
- **newspaper**: per capita Newspaper circulation (Banks)
- **tv**: per capita TVs (Banks)
- **radio**: per capita radios (Banks)
- **govpseat**: fraction of seats held by the government, which is calculated by dividing the number of government seats by total (government plus opposition plus non-aligned) seats (World Bank)
- **govpshare**: vote share of governing parties (World Bank)
- **prezahare**: vote share of the current president (World Bank)
- **ideology**: ideology of the government (left, right, center) (World Bank)

- `regpreced`: regional precedent of Human Rights Trials (cumulative years)
- `politydur`: polity duration (the years under the current democratic regime)

As I described above, the first step is to obtain the propensity scores for transitional justice measures. Typically, propensity scores are calculated using regression models for binary outcomes like logit or probit for the binary treatment. If we are interested only in the effects of human rights trials, we could calculate propensity scores for human rights trials, given covariates ( $Y_{x=1}|X$ ), with a binary logistic regression models and use them for further analyses (Guo and Fraser 2010; Holmes 2014). The main treatment variable of interest in our case is not binary, but is composed of multiple levels (eight). And each level represents a different combination of transitional justice, and we need to get propensity scores for each combination.

I use the pooled multinomial logit for extracting the propensity scores for each combination of transitional justice adoption. One of the crucial assumptions in causal inference is the conditional ignorability assumption, and the best way to avoid the violation of it is to include as many covariates as possible, although it is almost unattainable in observational settings. However, we could not expand the number of variables infinitely due to the degrees of freedom. Table 4.4 shows the ordered regression results with the inverse probability weights included. The description of resulting weight is as follows: mean: 8.344, media: 2.225, Min: 1.027, max: 276.2772. Table 4.4 shows the ordinal regression results with IPTW (with clustered standard errors for democratic regime). Figure 4.2 shows the density plots for the inverse probability weights for each treatment levels.


Figure 4.2: Density Plots of Stabilized Weights for Each Treatment Level



The next step is to include the censoring weight, which can be calculated by fitting the pooled logistic regression to a binary variable, censored or not. By running the logistic regression, we get the association between the covariates and the possibility of being censored, e.g., collapse of democratic regime. By subtracting propensity scores for censoring from 1, we can get the propensity scores for being uncensored, i.e., staying in the sample. Then, the scores are multiplied by the propensities in the previous years, unless the observation is censored in the given year. The distributions of censoring weight is shown in Figure 4.3.

The final weight is calculated by multiplying the simple weights by the censoring weights. The distribution of the final weights, stabilized censoring weight, is

Figure 4.3: Density Plots for Stabilized, Censoring, and Stabilized Censoring Weights



shown in Figure 4.3.

Table 4.4 shows the results of ordinal regression for physical integrity rights with and without weights. By including inverse probability weights as sampling weights in the model, we get the starkly different results.

Notable is that the signs of the coefficients for `growthrate`, `politydur` are changed in the model with weights, although many of them are insignificant. The coefficients for many variables, including `prezshare`, `govpshare`, `regpreced`, `lngdp`, and `newspapers` lose the significance at any level in the model with weights. The coefficients for `govpseat` and `politydur` gain the significance at varying levels.

Table 4.4: Ordered Regression Result 1

	Unadjusted				Adjusted			
	Robust				Robust			
	Coef.	Std. Err.	$z$	$p$	Coef.	Std. Err.	$z$	$p$
TJ level	.0652025	.0659902	0.99	0.323	.2963583	.1419195	2.09	0.037
Amnestylag	1.858309	.2019274	9.20	0.000	1.116663	.41364	2.70	0.007
polity	.1028604	.0677183	1.52	0.129	.1103989	.093958	1.17	0.240
govpseat	.0617749	.5350099	0.12	0.908	2.673744	1.450555	1.84	0.065
prezshare	.0162272	.0079223	2.05	0.041	.0124558	.0182305	0.68	0.494
ideology	-.1092452	.1150378	-0.95	0.342	-.0250748	.2531288	-0.10	0.921
govpshare	-.0094279	.0036522	-2.58	0.010	-.0121514	.0100722	-1.21	0.228
treaties	.0361642	.1382998	0.26	0.794	.3352976	.3825993	0.88	0.381
NGOs	-.0072133	.0952325	-0.08	0.940	-.2393307	.309671	-0.77	0.440
growthrate	.0007882	.037201	0.02	0.983	-.0386496	.0526596	-0.73	0.463
regprecd	.0100854	.0059367	1.70	0.089	.0212254	.0191167	1.11	0.267
lngdp	-.3489806	.1563788	-2.23	0.026	-.4632777	.3669481	-1.26	0.207
newspapers	.0310223	.0160574	1.93	0.053	.0054562	.0367945	0.15	0.882
radios	.002886	.0076856	0.38	0.707	.0336591	.0178962	1.88	0.060
lnpop	-.4219623	.1303624	-3.24	0.001	-.570637	.2304285	-2.48	0.013
popgrowth	-.3331213	.1883704	-1.77	0.077	-.4182222	.2259018	-1.85	0.064
politydur	.0019026	.0195552	0.10	0.922	-.0774569	.0408884	-1.89	0.058
cut1	-8.318268	3.52465			-11.5369	5.42377		
cut2	-5.168945	3.371914			-7.696834	5.187852		
cut3	-2.044685	3.393413			-5.550453	5.12735		
cut4	1.285413	3.410164			-1.267589	5.014871		
Log-likelihood	-486.6493				-2607.6869			
Pseudo- $R^2$	0.283				0.250			
Chi-squared	413.31				222.71			
$N$	522 (39 clusters)				522 (39 clusters)			

Note: The coefficients for region dummies are not reported.

The lag value of physical integrity rights is, as expected, the most powerful predictor for the current physical integrity rights with or without weighting. The results show that most of the variables lose the significance, although the pseudo- $R^2$  increases.

Based on these ordinal regression models, Tables 4.5 and 4.6 show the difference between the two regressions in predicted probabilities on each outcome. For the

calculation, transitional justice is (hypothetically) moving from none to all (0 to 8), but other values are set at the mean, except physical integrity lag, which is set at 1 (the lowest score).<sup>12</sup> The predicted probabilities indicate the probabilities that a democratic regime whose values are at the mean, except the past physical integrity (worst) is moving to each level of physical integrity. Not surprisingly, the predicted probabilities for  $y = 2$  are highest in every combination of transitional justice measure in the model without the weight. However, the model with weights show the different pictures, in which more nuanced interpretation can be possible. Overall, the predicted probabilities for  $Y = 3$  are highest in all combinations, except the all categories, i.e., the regime that adopted trials, truth commissions, and amnesties). The predicted probabilities for  $Y = 4$  are infinitesimally small in the model without weight, while those in the model with weights are noticeably higher. The new model with weights clearly show the varying effect of transitional justice measures on physical integrity.

Table 4.5: Predicted Probabilities for Select Scenario (unadjusted)

	None	Amnesties	TRCs	Trials	TRCs, Amnesties	Trials, Amnesties	Trials, TRCs	All
Pr( $y=1$ )	0.2022	0.1919	0.1820	0.1725	0.1634	0.1546	0.1463	0.1384
Pr( $y=2$ )	0.6531	0.6551	0.6564	0.6569	0.6566	0.6555	0.6536	0.6509
Pr( $y=3$ )	0.1373	0.1451	0.1532	0.1617	0.1705	0.1797	0.1892	0.1992
Pr( $y=4$ )	0.0071	0.0076	0.0081	0.0086	0.0092	0.0098	0.0105	0.0112
Pr( $y=5$ )	0.0003	0.0003	0.0003	0.0003	0.0003	0.0004	0.0004	0.0004

Figure 4.4a and 4.4b illustrate that the inclusion of the weights changes the predicted probabilities, especially those for  $Y = 3$  and  $Y = 4$  as the transitional

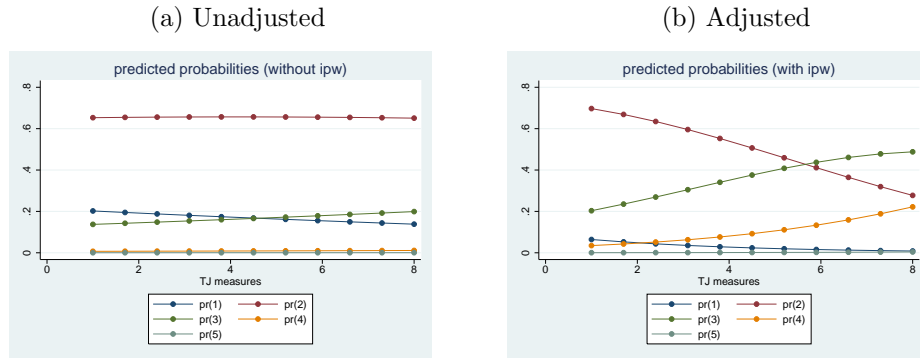
<sup>12</sup>Amnesty lag:1, polity: 6.878, govshare: 0.58, prezshare: 48.28, ideology: 4.70, govshare: 35.50, treaties: 2.51, NGOs: 2.15, growthrate: 2.42, regprecede: 22.88, lngdp: 7.44, newspapers: 7.41, radios: 32.17, lnpop: 16.48, popgrowth: 1.92, politydur: 10.41.



Table 4.6: Predicted Probabilities for Select Scenario (adjusted)

	None	Amnesties	TRCs	Trials	TRCs, Amnesties	Trials, Amnesties	Trials, TRCs	All
Pr(y=1)	0.0053	0.0040	0.0029	0.0022	0.0016	0.0012	0.0009	0.0007
Pr(y=2)	0.1940	0.1522	0.1180	0.0906	0.0691	0.0523	0.0395	0.0297
Pr(y=3)	0.4811	0.4567	0.4197	0.3739	0.3235	0.2725	0.2242	0.1807
Pr(y=4)	0.3131	0.3785	0.4477	0.5177	0.5850	0.0462	0.6985	0.7399
Pr(y=5)	0.0064	0.0086	0.0116	0.0155	0.0208	0.0277	0.0369	0.0491

Figure 4.4: Predicted Probabilities



justice measures vary from none to all, while the predicted probabilities for  $Y = 2$  decrease. The results show that the probabilities for improved physical integrity increase as the regime adopt more transitional justice measures.

## Discussion

In this chapter, I introduced the propensity score, the inverse probability treatment weight, and the inverse probability censoring weight and applied them to estimate the causal effects of transitional justice on physical integrity rights. I extended the inverse probability weights to the cases of multiple treatments (transitional justice) and right-censoring (the collapse of democratic regime). The purposes of creating those weights in this chapter are mainly three-fold: 1) to deal with the

combined treatments, instead of a single treatment, 2) to deal with the complex situations when the treatment is multi-level, instead of binary, and 3) to address the problem of censoring. Although I use the inverse probability weight to deal with complex treatments, i.e., multiple combinations of treatments, in this chapter, there is no reason not to extend the idea of inverse probability weight to adjust for various sorts of confounding such as time and geography. In the next chapter, I expand the idea of inverse probability weight to deal with time-varying confounding. Although time is one of the important concerns in statistical methodology, time-varying confounding has not been paid serious attention.

Using the adjustments by inverse probability weights, I showed that the estimates for the effects of transitional justice measures could be vastly different. More specifically, the advances are made on existing studies in two aspects: 1) in understanding transitional justice as a combination of multiple transitional justice measures, and 2) in incorporating the probability of regime collapse to the model. Using the procedure, called “doubly robust” procedure, I showed there are significant discrepancies in the estimates between the model with and without proper adjustment, and that the adjustment is a necessary step to deal with the selection problem in estimating the effect of transitional justice.

## Chapter 5

# Causal Inference for Dynamic Treatment Data

Although an increasing number of works have (re)appeared on causal analysis in recent decades, partly thanks to the potential outcomes framework,<sup>1</sup> they have largely been limited to static setting. The quantitative literature on policy and program evaluations are typically posed in a static single-shot treatment framework and hence the quantity of interest is the effect of the single-shot treatment on the single outcome (Bingham and Felbinger 2002). Timing and history of treatment(s) are usually ignored in this framework.<sup>2</sup>

The single-shot framework of causal analysis and its applications provide enormously useful (new) tools, but its applications are inevitably limited to the circumstances where the treatment occurs only once. This prototype, however, does not concur with many situations in which treatments occur multiple times over periods.<sup>3</sup>

---

<sup>1</sup>Although the potential outcomes framework revived the interest on causality and causal analysis in statistics, causality is still a concept that most statisticians avoid (Pearl 2000, 2009; Hoover 2001; Hernán et al. 2006).

<sup>2</sup>There are some ways to deal with the timing of the treatment in single-shot treatment framework (Hirano and Imbens 2004), but the question is how the timing of the treatment affects the treatments in the ensuing periods.

<sup>3</sup>The idea of dynamic treatment originated in epidemiology, and the well-known clinical examples include the treatments of AIDS, cancer and mental depression, which usually require a series of multiple treatments contingent on the development of the patient. Typically, the treatments include varying doses of the same medicine or a combination of different medicines/surgeries.

In these situations, target problems and treatments are usually continuously interacting among the covariates, the treatments and the (intermediate) outcomes, all of which are likely to be time-varying. Thus, the decision for treatments is rarely an once-and-for-all decision, but a sequential one in the sense that treatments are adjusted, changed, added or even discontinued (and replaced) based on ongoing progress, e.g., unexpected side effects, changed circumstances, or non-compliance. Treatment decisions are adjusted over time based on accruing observations on the subjects over time. The decision-maker is a *smart* player. Thus, the treatment needs to be considered as a dynamic sequence rather than as a static single-shot.

The dynamic treatment regime, or adaptive treatment strategy, is a function which takes in treatments and covariate history as arguments and outputs an action to be taken, i.e., a list of decision rules for allocation of treatments over time (Robins 1986; Murphy 2003). The treatment in each interval between measurements is not known at the start of a treatment regime (or treatment set), since treatment depends on subsequent time-varying variables that may be influenced by earlier treatment(s). Hence, the problem of finding the optimal treatment set is one of sequential decision-making, where an action that appears optimal in the short-term may not be a component of the optimal treatment in the longer-term. A treatment set is defined as optimal if it maximizes the mean response *at the end of the final time interval*. In this sense, adaptive treatment can be thought of as an algorithm that dictates how treatment of a problem should proceed over time (Chow and Chang 2006; Berry et al. 2010; Chen and Peace 2011).<sup>4</sup>

---

<sup>4</sup>The question of optimal sequencing, i.e., finding the optimal route to reach from point  $A$  to

The purpose of this chapter is to introduce the statistical tools to deal with time-varying confounding, to construct a causal framework to analyze the long term effect of transitional justice from the potential outcomes framework, and to estimate the long-term causal effect of human rights trials on human rights. In typical quantitative studies on transitional justice, transitional justice, as a whole or part of it, is considered a single shot treatment, whether and to what extent a government or a country adopted a transitional justice measure, and the quantity of interest is the effect of the adopted measure(s) on some desiderata such as human rights or stability. Typical examples in this line of research include the estimation of the effects of human rights trials or truth commissions on human rights, where human rights trials or truth commissions are considered single-shot treatment(s) (Sikkink and Walling 2007; Kim and Sikkink 2010; Olsen, Payne and Reiter 2010). Although this strategy can be effective in estimating the short-term effect of the specific transitional measures, it does not take into account the fact that the governments change, adjust, or abandon treatments over time in accordance with the changes in the covariates and intermediately observed outcomes. Hence, the estimates of the effectiveness of a certain measure(s) could be different depending on the time frame under discussion. The problem persists even in the longitudinal studies.

---

point  $B$ , is not new and has been traditionally studied in dynamic programming in computer science and optimal transport literature in mathematics to solve the sequential decision problems (Villani 2003, 2009). In the dynamic treatment context, however, dynamic programming requires modeling the longitudinal distribution of all covariates, intermediate and final outcomes (Bellman 1957; Powell 2007; Bertsekas 2007). However, the knowledge needed for dynamic programming is often unavailable and, by mis-specifying the distribution, treatment may be incorrectly recommended when no treatment effect actually exists. In this situation, the non-parametric framework could be useful. Another possibility is to use propensity score matching for each sequence of treatments (Lechner 2004, 2006).

The chapter proceeds as follows. The chapter begins with the description of methodological problems and notations for dynamic treatment regimes. Next, I discuss marginal structural modeling, proposed by Robins to deal with time-varying confounding. Then, the developed estimation strategies are applied in estimating the transitional justice's effects.

## The Methodological Framework to Estimate Dynamic Treatments

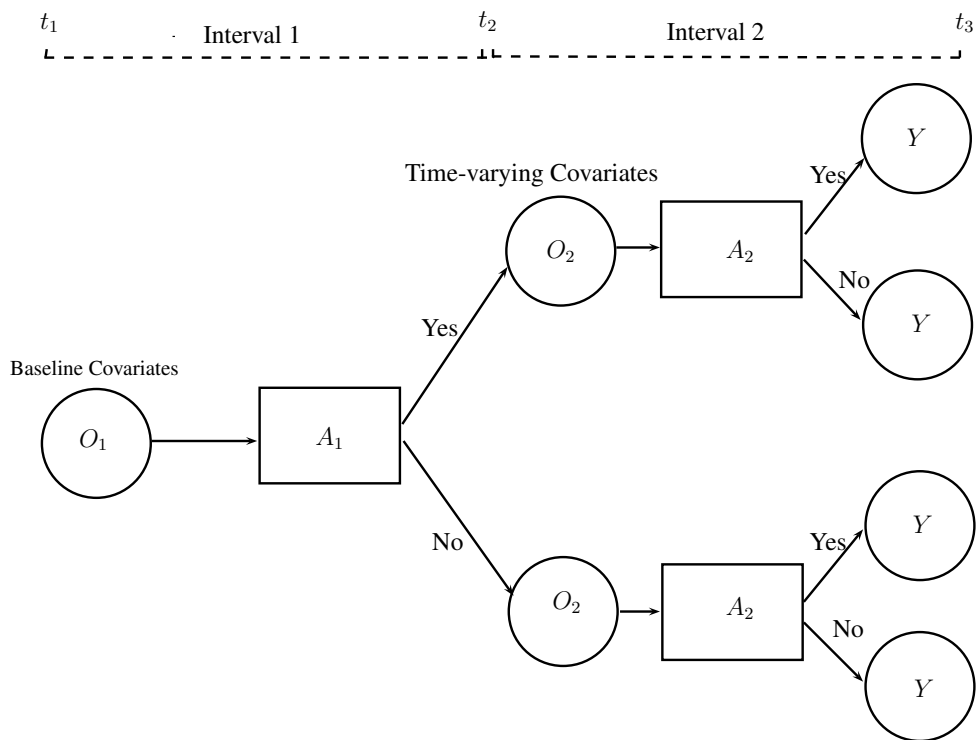
### Set-up and Notations

Treatments are given at  $K$  fixed times,  $t_1, t_2, \dots, t_K$ .  $O_j$  are the covariates measured prior to treatment at the beginning of the  $j$ th interval, while  $A_j$  is the treatment subsequent to having measured  $O_j$ .  $Y$  is the outcome observed at the end of interval  $K$ ; larger values of  $Y$  are considered favorable outcome for convenience. Random variables are upper case; specific values, or fixed functions are lower case. Denote a variable  $X_j$  at  $t_j$  and its history  $(X_1, X_2, \dots, X_j)$ , by  $\bar{X}_j$ . Finally, denote a treatment decision at  $t_j$  that depends on history,  $\bar{O}_j, \bar{A}_{j-1}$ , by  $D_j \equiv d_j(\bar{O}_j, \bar{A}_{j-1})$ .

For the purpose of illustration, let us begin with the simplest possible model with dynamic components. In the two-interval model, treatments are taken at two fixed times,  $t_1$  and  $t_2$ , with outcome measured at  $t_3$ .  $O_1$  and  $O_2$  are the covariates measured prior to the treatment of the first and the second intervals, respectively.  $O_1$  represents baseline covariates and  $O_2$  includes time-varying covariates that may depend on treatment received at  $t_1$ . The treatment given subsequent to observing  $O_j$  is  $A_j$ ,  $j = 1, 2$ . Response  $Y$  is measured at  $t_3$ . Thus, the chronological order of

the variables is  $(O_1, A_1, O_2, A_2, Y)$ , as Figure 5.1 shows.

Figure 5.1: The Chronological Order of the Data in Two Interval Model



The two interval model in Figure 5.1 can be re-framed in a potential outcomes framework. The quantity of interest in this framework is the value of a covariate or the outcome that would result if a subject were assigned to a counterfactual treatment, which would be different from the actual treatment. In the two-interval case we denote by  $O_2(a_1)$  a subject's potential covariate at the beginning of the second interval if treatment  $a_1$  is taken by that subject, and  $Y(a_1, a_2)$  denotes the potential end-of-study outcome if regime  $(a_1, a_2)$  is followed.

Potential outcomes adhere to the axiom of consistency:  $O_2(a_1) \equiv O_2$ , wher-

ever treatment  $a_1$  is actually received and  $Y(a_1, a_2) \equiv Y$  wherever  $(a_1, a_2)$  is received. In words, the actual and the counterfactual covariates (or outcome) are equal when the potential regime is the regime actually received.<sup>5</sup>

## Main Assumptions

To re-iterate from Chapter 2, two assumptions are required to estimate the effects of adaptive treatments: stable unit treatment value assumption (SUTVA) and sequential ignorability (Assumptions 1, 2, 3 and 4).

1. Stable Unit Treatment Value Assumption (SUTVA): A subject's outcome is not influenced by other subjects' treatment allocation (Rubin 2006).
2. Sequential Ignorability (no unmeasured confounders) (Robins 2004) For any regime  $\hat{a}_K$ ,

$$A_j \perp\!\!\!\perp (O_{j+1}(\bar{a}_j), \dots, O_K(\bar{a}_{K-1}), Y(\bar{a}_K)) \mid \bar{O}_j, \bar{A}_{j-1} = \bar{a}_{j-1}.$$


In principle, the second assumption always holds if subjects are sequentially randomized. In two intervals for any  $a_1, a_2$  we have  $A_1 \perp\!\!\!\perp (O_2(a_1), Y(a_1, a_2)) \mid O_1$  and  $A_2 \perp\!\!\!\perp Y(a_1, a_2) \mid (O_1, A_1 = a_1, O_2)$ , that is, conditional on treatment history, treatment received in any interval is independent of any future potential outcome.

---

<sup>5</sup>According to Robins, an estimator  $\hat{\psi}_n$  of the parameter  $\psi^\dagger$  is  $\sqrt{n}$ -consistent if  $\sqrt{n}(\hat{\psi}_n - \psi^\dagger) = O_p(1)$  with respect to any distribution  $P$  in the model. Similarly, if under  $P$ ,  $\sqrt{n}(\hat{\psi}_n - \psi^\dagger) \xrightarrow{D} T$ , and  $E|T| < \infty$ , then  $E(T)$  is the  $(\sqrt{n})$  asymptotic bias of  $\hat{\psi}_n$  under  $P$ , which we can denote  $AsyBias_p(\hat{\psi}_n)$ . If for some law  $P$  in the model,  $AsyBias_p(\hat{\psi}_n) = 0$ , then  $\hat{\psi}_n$  is said to be asymptotically unbiased under  $P$ .  $g$ -estimators are  $\sqrt{n}$ -consistent under all laws, however they are not asymptotically unbiased under certain distributions of the data, which Robins call "exceptional laws."



Figure 5.2: Directed Acyclic Graphs Representing Different Assumptions about Sequential Ignorability (Blackwell 2012)



The assumption of sequential ignorability extends the conditional ignorability assumption to time-varying actions. It states that action decision at time  $t$  is independent of the potential outcomes, conditional on the covariate and action histories up to that point. Figure 5.2a and 5.2b show the situations where sequential ignorability holds and fails (due to confounding from  $U$ ), respectively. If decisions are made randomly, then clearly this assumption holds (Blackwell 2012).

We can see from these assumptions that the estimation is possible only within the feasible set of treatment regimes, which is sometimes called ‘positivity’ assumption in the literature. In other words, if none of the subjects are actually in a possible set, the set is not used in estimation. In the two interval model, four combinations of treatments in each period are possible: (non-treatment, non-treatment), (non-treatment, treatment), (treatment, non-treatment), and (treatment, treatment). Figure 5.1 shows the chronological order of the data for this scenario.

If we are forced to consider the treatment as a single-shot in a single period, all of three treatment sequences except the first one, i.e., (non-treatment, non-treatment), are considered same. A more serious problem arises regarding the second and the third combination. Intermediate outcomes, measured at  $t_2$ , will be pre-

treatment variables for the (non-treatment, treatment) group, and post-treatment variables for the (treatment, non-treatment) group. The dynamic application fits, at best, awkwardly into the single-shot framework.

## **Modeling Dynamic Treatments with Marginal Modeling**

There are two fundamentally different ways of handling dynamic treatment and time-varying covariates within the potential outcomes framework (Aalen, Borgan and Gjessing 2008). One is the marginal procedure, where the aim is to estimate a treatment effect, and all other effects are considered nuisances to be adjusted for. Elaborate weighting procedures have been developed during the process. These procedures also focus mainly on estimating what would happen in a population with the same composition as the available sample, thereby limiting the generalizability. The other is joint modeling, where all or most components in the process are modeled. For example, the change of treatment is explicitly modeled as a function of previous observations, as are also time-varying covariates. The primary focus in this chapter is more on the first set of the techniques because the state of knowledge necessary for joint modeling is not yet reached in political science in general, comparative politics in particular.

### **Marginal Structural Models and Inverse Probability Weighting**

The aim of the marginal structural model, which was first introduced by Robins (1986), is to estimate the effect of a treatment in the presence of time-varying covariates, which may both influence and be influenced by the treatment.

Essentially, the approach consists in weighting the observations appropriately before estimation.<sup>6</sup>

For illustration, let  $A(t)$  be 1 if the subject is on a treatment and 0 otherwise, and let  $\bar{A}(t)$  denote the individual's treatment history up to, but not including,  $t$ . Furthermore, let  $\bar{O}(t)$  denote the corresponding history of time-varying covariates, and let  $V$  denote the baseline values of the covariates. It is common to define the weights based on discretization of the time interval. Usually the weights are determined by logistic regression, so the interval should not be too small. Let  $t$  be the end of time interval. The following weights are used, where multiplication is performed over all time intervals up to time  $t$ .

$$sw_i(t) = \prod^{\text{int}(t)} \frac{\Pr(A(k) = a_i(k) | \bar{A}(k-1) = \bar{a}_i(k-1), V = v)}{\Pr(A(k) = a_i(k) | \bar{A}(k-1) = \bar{a}_i(k-1), \bar{O}(k-1) = \bar{o}_i(k-1))}. \quad (5.1)$$

Then, the final step is to carry out standard analysis, e.g. regression analysis, where each observation in the risk set is weighted according to the above weights,

---

<sup>6</sup>Take the example of zidovudine treatment on HIV patients. Robins and co-authors study the effect of the medication zidovudine on the survival of HIV-positive patients. A number of time-varying covariates are measured, including CD4 count. CD4 is a time-varying confounder because it is a risk factor for mortality and also a predictor of initiation of zidovudine. Here, it is assumed that an individual stays on zidovudine therapy once it is started. When just treatment is included as covariate, i.e., on zidovudine or not, the hazard ratio is estimated to be 3.55 (2.9-4.3), and when including baseline covariates, estimate is adjusted to 2.32 (1.9-2.8). Hence, one may get the impression that zidovudine increases the risk of dying. The reason, of course, is that treatment started when CD4 count is low, so one has to include information on the time-varying CD4 count. Doing that through a simple time-varying Cox model would not be sufficient because zidovudine therapy is influenced by and influences CD4 count. Marginal structural model solves this problem by weighting the observations in an appropriate way. In the presence of confounding, the stabilized weights increase the efficiency. When carrying out the weighted analysis, the hazard ratio for treatment is reduced to 0.74 (0.57 - 0.96), showing that there is no effect of treatment.

and where treatment as well as  $V$  are included as covariates.

These are inverse probability of treatment weights (IPTW) as I introduced in Chapter 4, but with time-varying confounding adjusted. To re-iterate, they are closely related to the propensity scores proposed by Rosenbaum and Rubin (1983) and by Ertefaie and Stephens (2010). As discussed in Chapter 4, the rationale for IPTW is that the observations are weighted according to the probability of observed treatment given their covariates. Those who have unusual treatments will be over-weighted, and those who have a common treatment will be under-weighted. Initially, the weights are the product formed from the denominators in Equation (5.1). However, it has been found that more efficient estimators are achieved by stabilizing the weights, that is, multiplying with the numerator terms in Equation (5.1). When there is no time-varying confounding, the numerators will equal the denominators and give unweighted estimation.

The stabilized version of inverse probability censoring weight is:

$$sw_i(t) = \prod^{\text{int}(t)} \frac{\Pr(A(k) = a_i(k) | \bar{A}(k-1) = \bar{a}_i(k-1), V = v, T > k)}{\Pr(A(k) = a_i(k) | \bar{A}(k-1) = \bar{a}_i(k-1), \bar{O}(k-1) = \bar{o}_i(k-1), T > k)}. \quad (5.2)$$

### Structural Nested Mean Model (SNMM) and $g$ -estimation

Another estimation strategy is called  $g$ -estimation, which is originally sketched by Robins and refined by Murphy in a series of papers in the biostatistics literature (Robins et al. 1992; Robins 2004; Robins, Orellana and Rotnitzky 2008; Murphy 2003; Moodie, Richardson and Stephens 2007). Although Robins proposed a number of estimating equations for finding optimal regimes using SNMMs, I focus on

a particular subclass of SNMMS, called *optimal blip-to-zero* functions, denoted by  $\gamma_j(\bar{o}_j, \bar{a}_j)$  and defined as the expected difference in outcome when using the ‘zero’ regime instead of treatment  $a_j$  at  $t_j$ , in subjects with treatment and covariate history  $\bar{o}_j, \bar{a}_{j-1}$  who subsequently receive the optimal regime. The ‘zero’ regime, which is denoted  $0_j$ , can be thought of as placebo or standard care in clinical trial or inaction in policy implementation (in our case, virtual immunity); the optimal strategy at time  $j$  is denoted as  $d_j^{opt}(\bar{o}_j, \bar{a}_{j-1})$ . The optimal treatment regime subsequent to being prescribed  $a_j$  (or  $0_j$ ) at  $t_j$  may depend on prior treatment and covariate history, that is, what is optimal subsequent to  $t_j$  may depend both on the treatment received at  $t_j$  and on  $(\bar{o}_j, \bar{a}_{j-1})$ . In the two-interval case, the blip functions are:

$$\begin{aligned}\gamma_1(o_1, a_1) &= E[Y(a_1, d_2^{opt}(o_1, a_1, O_2(a_1))) - Y(0_1, d_2^{opt}(o_1, 0_1, O_2(0_1))) | O_1 = o_1], \\ \gamma_2(\bar{o}_2, \bar{a}_2) &= E[Y(a_1, a_2) - Y(a_1, 0_2) | (\bar{O}_2, A_1) = (\bar{o}_2, a_1)].\end{aligned}$$

At the last interval (the second in a two interval model), there are no subsequent treatments, so the blip  $\gamma_2(\cdot)$  is simply the expected difference in outcomes for subjects having taken treatment  $a_2$  as compared with the zero regime,  $0_2$ , among subjects with history  $\bar{o}_2, a_1$ . At the first interval, the blip  $\gamma_1(\cdot)$  is the expected (conditional) difference between the counterfactual outcome if treatment  $a_1$  was given in the first interval and optimal treatment was given in the second and the counterfactual outcome if the zero regime was given in the first interval and optimal treatment was given in the second interval.

To use  $\psi$  to denote the parameters of the optimal blip function model, and

$\psi^\dagger$  to denote the true values, for example, a linear blip  $\gamma_j(\bar{o}_j, \bar{a}_j; \psi) = a_j(\psi_{j0} + \psi_{j1}o_j + \psi_{j2}o_j^2 + \psi_{j3}a_{j-1})$  implies that, conditional on prior treatment and covariate history, the expected effect of treatment  $a_j$  on outcome, provided optimal treatment is subsequently given, is quadratic in the covariate  $o_j$  and linear in the treatment received in the previous interval. Note that in this example, the blip function is a simple function of the covariate multiplied by the treatment indicator  $a_j$ .<sup>7</sup>

Robins proposes finding the parameters  $\psi$  of the optimal blip-to-zero function via  $g$ -estimation. For two intervals, let

$$H_1(\psi) = Y - \gamma_1(O_1, A_1; \psi) + [\gamma_2(\bar{O}_2, (A_1, d_2^{opt}(O_1, A_1, O_2(A_1))); \psi) - \gamma_2(\bar{O}_2, \bar{A}_2; \psi)],$$

$$H_2(\psi) = Y - \gamma_2(\bar{O}_2, \bar{A}_2; \psi).$$

$H_1(\psi)$  and  $H_2(\psi)$  are equal in expectation, conditional on prior treatment and covariate history, to the potential outcomes  $Y(0_1, d_2^{opt}(O_1, 0_1, O_2(0_1)))$  and  $Y(A_1, 0_2)$ , respectively. For the purpose of constructing an estimating procedure, we must specify a function  $S_j(a_j) = s_j(a_j, \bar{O}_j, \bar{A}_{j-1}) \in \mathbb{R}^{dim(\psi)}$  which depends on variables

---

<sup>7</sup>Non-linear SNMMs are possible and may be preferable for continuous doses of treatment. For example, the SNMMs corresponding to the models in Murphy for both continuous and binary treatments are quadratic functions of treatment. It is equally possible to specify SNMMs where parameters are common across intervals. For instance, in a two interval setting, the following blips may be specified:  $\gamma_1(o_1, a_1) = a_1(\psi_0 + \psi_1 o_1)$  and  $\gamma_2(\bar{o}_2, \bar{a}_2) = a_2(\psi_0 + \psi_1 o_2)$ . In this example, the same parameters,  $\psi_0$  and  $\psi_1$  are used in each interval and thus are said to be shared. The implied optimal decision rules are  $I[\psi_0 + \psi_1 o_j > 0]$  for intervals  $j = 1, 2$ . Practically, sharing of parameters by blip functions/decision rules would be appropriate if the researcher believes that the decision rule in each interval  $j$  is the same function of the covariate  $o_j$  (and thus does not depend on the more distant past). Note that if  $a_j$  takes multiple levels then we are free to allow the blip function to take different functional forms for different values of  $a_j$ .

that are thought to interact with treatment to influence outcome. For example, if the optimal blip at the second interval is linear, say,  $\gamma_2(\bar{o}_2, \bar{a}_2) = a_2(\psi_0 + \psi_1 o_2 + \psi_2 a_1 + \psi_3 o_2 a_1)$ , a common choice for this function is  $S_2(a_2) = \frac{\partial}{\partial \psi_2} \gamma_2(\bar{O}_2(A_1, a_2)) = a_2 \cdot (1, O_2, A_1, O_2 A_1)^\top$  as the blip suggests that the effect of the treatment on outcome at  $t_2$  is influenced by the covariates at the start of the second interval and by treatment at  $t_1$ . Let

$$U_j(\psi) = \{H_j(\psi) - E[H(\psi)|\bar{O}_j, \bar{A}_{j-1}]\}\{S_j(A_j) - E[S_j(A_j)|\bar{O}_j, \bar{A}_{j-1}]\} \quad (5.3)$$

In most of the distributions, if  $U(\psi) = \sum_{j=1}^2 U_j(\psi)$ , then  $E[U(\psi^\dagger)] = 0$  is an unbiased estimating equation from which consistent estimates  $\hat{\psi}$  of  $\psi^\dagger$  may be found. Robins proves that estimates found by solving Equation (5.3) are consistent if the expected counterfactual model,  $E(H_j(\psi)|\bar{O}_j, \bar{A}_{j-1})$  is correctly specified. Since, for consistency, only one of the models need be correct, this procedure is said to be doubly robust. At exceptional laws the estimators are consistent but not asymptotically normal and not asymptotically unbiased. At non-exceptional laws the estimators are asymptotically normal under standard regularity conditions but are not in general efficient without a special choice of the function  $S_j(A_j)$ .

A less efficient, singly robust version of Equation (5.3) simply omits the expected counterfactual model:

$$U_j^*(\psi) = H_j(\psi)\{S_j(A_j) - E[S_j(A_j)|\bar{O}_j, \bar{A}_{j-1}]\} \quad (5.4)$$

Estimates found via Equation (5.4) are consistent if the model for treatment allocation,  $\Pr_j(a_j|\bar{O}_j, \bar{A}_{j-1})$  is correctly specified.

Exact solutions to Equation (5.3) and Equation (5.4) can be found when optimal blips are linear in  $\psi$  and the parameters are not shared between intervals. An algorithm for solving the doubly robust estimating equation Equation (5.3) in the two-interval case is as follows, using  $\mathbb{P}_n$  to denote the empirical average operator.

1. Estimate the nuisance parameters of the treatment model at time 2; that is, estimate  $\alpha_2$  in  $p_2(a_2|\bar{O}_2, A_1; \alpha_2)$ .
2. Assume a linear model for the expected counterfactual,  $E[H_2(\psi_2)|\bar{O}_2, A_1; \zeta_2]$ . Express the least squares estimate  $\hat{\zeta}_2(\psi_2)$  of the nuisance parameter  $\zeta$ , explicitly as a function of the data and the unknown parameter,  $\psi_2$ .
3. To find  $\hat{\psi}_2$ , solve  $\mathbb{P}_n U_2(\psi_2) = 0$ ; that is, solve

$$\mathbb{P}_n \{H_2(\psi_2) - E[H_2(\psi_2)|\bar{O}_2, A_1; \hat{\zeta}_2(\psi_2)]\} \{S_2(A_2) - E[S_2(A_2)|\bar{O}_2, A_1; \hat{\alpha}_2]\} = 0.$$

4. Estimate the nuisance parameters of the treatment model at time 1; that is, estimate  $\alpha_1$  in  $p_1(a_1|O_1; \alpha_1)$ . Plug  $\hat{\psi}_2$  into  $H_1(\psi_1, \psi_2)$  so that only  $\psi_1$  is known.
5. Assuming a linear model for  $E[H_1(\psi_1, \hat{\psi}_2)|O_1; \zeta_1]$ , the least squares estimate  $\hat{\zeta}_1(\psi_1, \hat{\psi}_2)$  of  $\zeta_1$  can again be expressed directly in terms of  $\psi_1$ ,  $\hat{\psi}_2$  and the data.
6. Solve  $\mathbb{P}_n U_1(\psi_1, \hat{\psi}_2) = 0$  to find  $\hat{\psi}_1$ ; that is, solve

$$\mathbb{P}_n \{H_1(\psi_1, \hat{\psi}_2) - E[H_1(\psi_1, \hat{\psi}_2)|O_1; \hat{\zeta}_1(\psi_1, \hat{\psi}_2)]\} \{S_1(A_1) - E[S_1(A_1)|O_1; \hat{\alpha}_1]\} = 0$$



## Transitional Justice as Dynamic Treatment Regime

The question is now whether and how transitional justice can be modeled as a dynamic treatment regime. Using the framework of dynamic treatment regime, the idea can be stated as *the democratic government adopts transitional justice measures subsequent to observing the human rights situations. And they continue, adjust, or stop those measures as they observe the human rights situations in the ensuing periods.* Most of the small but growing literature on transitional justice tends to consider various transitional justice measures as separate and single-shot treatments and attempts to estimate the causal effect of those measures, typically human rights trials, within a single time frame (Sikkink and Walling 2007; Olsen, Payne and Reiter 2010). The research questions take the form of whether human rights trials were effective in improving human rights conditions in three years by comparing the mean human rights indicators of the countries that adopted trials (treated) and those that did not.

The emerging questions in applying potential outcomes framework are whether the two main assumptions of potential outcomes framework (SUTVA and sequential ignorability) are not violated. It is almost a truism in political science that we need to consider the neighborhood effect (diffusion) or learning effect, especially when we study policy implementation. For example, truth commissions, one of the most widely known transitional justice measures, have been widely and newly employed across the world since the much-heralded South African Truth and Reconciliation Commission in 1994-5. Dissimilar to clinical studies, where SUTVA is rarely violated, this poses a serious problem in social science, as I discussed in Chapter 2.

The same question can also be raised about sequential ignorability.<sup>8</sup> Unlike clinical trials or experimental setting, where sequential randomization guarantees sequential ignorability as in SMART design (sequential, multiple assignment, randomized trials), the assumption is not easy to retain in observation settings, albeit we attempt to include many variables in order to minimize unmeasured confounding.

Both of these are serious concerns for most of social science applications of potential outcomes framework, but, unfortunately, we do not have ready-to-use solutions to these problems when the assumptions are believed to be violated.

### **Application of Inverse Probability Weights for Adjustment: Physical Integrity**

In this section, I extend the discussion of the inverse probability weights for the combined transitional justice measures in Chapter 4 to handle the multiple treatments in dynamic settings. Inverse probability weights for the treatment for the first interval are calculated by pooled logistic regression with baseline covariates. And the propensity scores for the second interval are calculated with the covariates at the covariates at  $t$ ,  $t - 3$ , and treatment history, i.e., the adopted transitional justice measures, in the last three years. The same procedure continues for the second (year 3 -5) and the third intervals (year 6-10). These elements constitute the denominator of the stabilizing weight. The numerator is calculated in the similar procedure, but only with baseline covariates and the treatment history, i.e., the adopted transitional justice measures up to that point.

---

<sup>8</sup>The same question is raised on ignorability assumption in mediation analysis.

Instead of including the trials, truth commissions, and amnesties separately or in combination, I created a treatment variable with multiple level for the first three years, the third-the fifth years, and the sixth-tenth year, which is composed of eight ordinal categories as I described in Chapter 4. Multinomial logistic regressions are used to calculate the propensity scores for treatment levels in each interval. A series of binary logit models are used for calculating the censoring weight for each interval, and the censoring weight for each observation is multiplied by  $(n - 1)$  observation if it is not censored. The final weight is the product of the stabilized weights and censoring weights. The distribution of the resulting weight is: min: 0.14, mean: 1.88, median: 1.22, max: 155.58. Figure 5.3 shows the density plots of the stabilized weights, censoring weights, and stabilized censoring weights.

Figure 5.3: Density Plots for the Weights




Table 5.1 compares the ordered logit regression results with clustered standard error (for the country) with and without stabilized weights. In addition to the

variables included in the models in Chapter 4, the following variables are included:  
*cumtrial*: cumulative years of human rights trials, *cumtrc*: cumulative years of truth commissions, *cumamnesty*: cumulative years of amnesties.

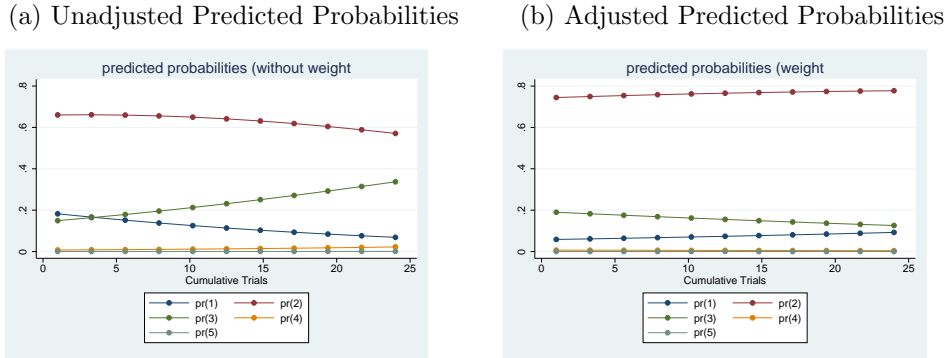
Table 5.1: Ordinal Regression Result 2

	Without Adjustment				With Adjustment			
	Coef.	Std. Err.	<i>z</i>	<i>p</i>	Coef.	Std. Err.	<i>z</i>	<i>p</i>
<i>Amnestylag</i>	1.80239	.2107807	8.55	0.000	1.981606	.3317394	5.97	0.000
<i>cumtrial</i>	.0480515	.0244435	1.97	0.049	-.0215058	.0746347	-0.29	0.773
<i>cumtrc</i>	.0810651	.0481436	1.68	0.092	.0092278	.1124441	0.08	0.935
<i>cumamnesty</i>	-.0378638	.0748975	-0.51	0.613	-.0424928	.1828073	-0.23	0.816
<i>polity</i>	.075327	.0745586	1.01	0.312	-.4018656	.2200477	-1.83	0.068
<i>ideology</i>	-.1276484	.115822	-1.10	0.270	-.2292793	.2361837	-0.97	0.332
<i>govpseat</i>	-.0314685	.5829559	-0.05	0.957	.9174179	1.544403	0.59	0.552
<i>prezshare</i>	.0163122	.0079573	2.05	0.040	.0223974	.0210402	1.06	0.287
<i>govpshare</i>	-.0092334	.0038688	-2.39	0.017	-.0203177	.0097801	-2.08	0.038
<i>treaties</i>	.1439451	.1620908	0.89	0.375	.2676362	.5600176	0.48	0.633
<i>NGOs</i>	-.0022971	.0984058	-0.02	0.981	-.1380134	.1509442	-0.91	0.361
<i>growthrate</i>	-.0056077	.0365477	-0.15	0.878	.1043914	.040928	2.55	0.011
<i>regpreced</i>	.0054947	.0076259	0.72	0.471	.056495	.0158902	3.56	0.000
<i>lngdp</i>	-.3824003	.1710121	-2.24	0.025	-.0248487	.4376871	-0.06	0.955
<i>newspapers</i>	.0336718	.017431	1.93	0.053	.0229005	.0364648	0.63	0.530
<i>radios</i>	.0028024	.0081898	0.34	0.732	.0378877	.0210279	1.80	0.072
<i>lnpop</i>	-.4605422	.1249696	-3.69	0.000	-1.275019	.2768758	-4.61	0.000
<i>popgrowth</i>	-.3789189	.2051027	-1.85	0.065	-.0047143	.6444313	-0.01	0.994
<i>politydur</i>	-.0146087	.0273982	-0.53	0.594	-.049244	.0488552	-1.01	0.313
<i>cut1</i>	-9.766096	3.554908			-22.14288	4.777119		
<i>cut2</i>	-6.586578	3.357268			-17.95975	4.449032		
<i>cut3</i>	-3.418677	3.377402			-14.41131	4.365147		
<i>cut4</i>	-.0904973	3.409142			-10.35703	4.378768		
Log-likelihood		-483.5289				-333.3439		
Pseudo- $R^2$		0.288				0.417		
$\chi^2$		280.12				472.73.92		
<i>N</i>		522				222		

Note: The coefficients for region dummies are not reported.

Figure 5.4a and 5.4b show the predicted probabilities for each outcome as the cumulative years of trials (hypothetically) move from 0 to 24 when physical integrity

Figure 5.4: Predicted Probabilities according to the Years in Trial



lag is 1 (the lowest score), and other values are set to the mean. Figure 5.4a shows that the probabilities for falling into level 3 are increasing as the trial years increase.

Noticeable change is only seen in the predicted probabilities for physical integrity score 3. In the model without weight, the probabilities for moving from 1 to 3 are increasing as cumulative years of trials increase, while they show slight decrease over time, while the probabilities for falling into level 1 decrease as trial years increase. In Figure 5.4b, the probabilities for falling into level 3 slightly decrease as trial years increase. Noteworthy is the probabilities for falling into level 2, which are higher in the model with weights.

**Application of *G*-estimation: Democratic Regime Survival**

As an illustration, I construct a simple scenario, where the outcome variable is the survival of the democratic regime. The treatment variable is human rights trials, and an additional explanatory variable is physical integrity rights. I used the logistic regression for predicting outcome variable *y*, and ordinal regression for predicting the

current value of physical integrity, and logistic regression for predicting the current treatment. The number of years under the current democratic regime is used for time variable. In summary, following variables and functions are used for this particular  $g$ -estimation:

Table 5.2: Functional Forms of  $g$ -estimation

Variable	function	Independent Variables
Regime Survival	logit	lag Physical Integrity, lagged Cumulative Trials
Physical Integrity	regression	lag Physical Integrity, lag Trials
Trials	logit	lag Physical Integrity, lag Trials

Note here that physical integrity is time-varying, and trials are also time-varying.  $G$ -estimation formula estimates for the parameters of the specified marginal structural model. The  $g$ -estimation estimate of coefficient is -0.3582, and the bootstrap standard error is 0.0367. For an easier interpretation of the results, Table 5.3 shows the cumulative incidence under each of the four treatment regime scenario. Here, the four treatment regimes (intervention) include: 1) trials in ten years, 2) trials only in the first three years, 3) trials only in the first five years after transition, and 4) no trials.<sup>9</sup> The results show the fourth scenario has the highest probability of regime collapsing (0.787), while the scenario 1 has the lowest. Note that the results in Table 5.3 is the estimates on the static regime.<sup>10</sup>

Next, I construct select dynamic regimes with the following scenarios.

---

<sup>9</sup>The treatment sequences in each regime are as follows: Regime 1: (1,1,1,1,1,1,1,1,1), Regime 2: (1,1,1,0,0,0,0,0,0), Regime 3: (1,1,1,1,1,0,0,0,0), Regime 4: (0,0,0,0,0,0,0,0,0).

<sup>10</sup>Part of the  $g$ -estimation in this section was performed by a Stata package, `gformula`, written by Rhian Daniel, Bianca De Stavola, and Simon Cousens.

Table 5.3: *G*-estimation Results: Static Regime

	G-computation estimate of cum. incidence	Bootstrap Std. Err.	z	p
Int. 1	.23	.020746	11.09	0.000
Int. 2	.396	.0275824	14.36	0.000
Int. 3	.28	.0255406	10.96	0.000
Int. 4	.787	.0241751	32.55	0.000
Obs. regime simulated	.502	.022168	22.65	0.000
observed	.519			

- Regime 1: Trials in the first ten years and the physical integrity is below 3, No trials in the first ten years if the physical integrity is 3 or above.
- Regime 2: No Trials after year 3 and the physical integrity is below 3, Trials in the first three years if physical integrity is less than 3.
- Regime 3: No trials after five years and the physical integrity is below 3, Trials in the first five years and physical integrity is lower than 3.
- Regime 4: No trials at all if the physical integrity is below 3.

The results in Table 5.4 show that none of the above regimes are effective in maintaining democratic stability.


Figure 5.5b shows the survival time for each intervention and the observed, which shows that none of these dynamic regimes are effective in maintaining democratic regime.

Table 5.4:  $G$ -estimation Results: Dynamic Regime


	G-computation estimate of cum. incidence	Bootstrap Std. Err.	z	p
Int. 1	1	.0105576	94.72	0.000
Int. 2	1	.0101482	98.54	0.000
Int. 3	1	.0088004	113.63	0.000
Int. 4	1	.0079606	125.62	0.000
Obs. regime simulated	1	.0078517	127.36	0.000
observed (o)	.7621951			
observed (l)	.5853659			

Figure 5.5: Kaplan-Meier Estimates for Regime Survival with Dynamic Regime

(a) Kaplan-Meier Estimates for the Data



(b) Kaplan-Meier Estimates for Simulated and Observed Regimes



## Discussion

In this chapter, I discussed the key components of dynamic treatment regime, marginal structural modeling, structural nested modeling, and various strategies to find the optimal dynamic treatment regime. To summarize, the primary purpose of all these is to find the best way to estimate the causal effects in the presence of confounding due to the time-varying treatments and covariates. I also described two estimation strategies for dynamic treatment regime: inverse probability weight



and  $g$ -estimation. The application of inverse probability weight is a continuation of the discussion in Chapter 4, but, the weight is used to deal with time-varying confounding in this chapter.  $G$ -estimation is more directly targeting to the provide the individualized subject-specific guideline for treatment strategies. We need to have clearer ideas on the functional forms among the treatments and covariates for structural nested models.

Although statistical framework and the estimation strategies that I introduced in this chapter could be quite useful, especially when such adjustment methods as matching and stratification are not feasible, they have rarely been discussed or used in political science. Only recently have the ideas of inverse probability weights begun to be introduced and applied in political science (Glynn and Quinn 2010; Blackwell 2012). In the next chapter, I discuss the reinforcement learning as conceptual links to dynamic treatment regimes that could provide the theoretical structure for alternative modeling of dynamic treatment regime.

## Chapter 6

# Dynamic Treatment Regime From Machine Learning Perspective

In the previous chapters, I have argued that constructing optimal dynamic treatment regime for chronic problems is a problem of multi-stage decision making about the best sequence of treatments. This problem bears strong resemblance to the problem of reinforcement learning in machine learning. Reinforcement learning is a branch of machine learning that deals with the multi-stage, sequential decision making by a learning agent (Sutton and Barto 1998). In this framework, a learning agent tries to optimize the total amount of reward it receives when interacting with an uncertain environment. Unlike supervised and unsupervised learning, this branch of machine learning is relatively less known in political science despite its obvious applicability in modeling complex strategic situations.<sup>1</sup> The purpose of this chapter is to introduce the basic concepts of reinforcement learning, to link them to the dynamic treatment regimes and the decision theory, and to develop a framework that will enable us to treat the problem of estimating optimal dynamic treatment regimes. Although part of this chapter is introduction to reinforcement learning,

---

<sup>1</sup>The applications of supervised and unsupervised learning techniques have been used in political science for text analysis. For some recent works that apply machine learning techniques to political science, see Quinn et al. (2010); Hopkins and King (2010); Grimmer (2013).

my purpose is not to provide a comprehensive treatment of machine learning in general or reinforcement learning in particular (see Sutton and Barto (1998); Buşoniu et al. (2010); Szepesvari (2010); Wiering and van Otterlo (2012); Schwartz (2014) for comprehensive and detailed discussions of the topic). My discussion in this chapter is confined to some aspects of reinforcement learning that can be applied to the dynamic treatment regime.

As is well known, the first serious attempt to solve multi-stage sequential decision problems is dynamic programming, which originates from Bellman (1957). However, classical dynamic programming algorithms are of limited practical utility in reinforcement learning because of the following two reasons. First, as I discussed in Chapter 5, they require the complete knowledge about the multivariate distribution of the data. In many areas of study, it is often impractical or impossible to assume full distributional knowledge of all variables. Second, dynamic programming methods are computationally very expensive, and they become hard to manage even in moderately high dimensional problems (curse of dimensionality). Although dynamic programming is still important as a theoretical foundation for reinforcement learning, the field of modern reinforcement learning experienced a major breakthrough when Watkins (1989) developed  $Q$ -learning ( $Q$  for “quality”), a method to solve multi-stage decision problems based sample data trajectories. For this reason, the following discussions of the reinforcement learning will be largely restricted to  $Q$ -learning and its applications to the dynamic treatment regime.<sup>2</sup>

---

<sup>2</sup>The other modeling strategies for optimal treatment regime from reinforcement learning perspective could include  $A$ -learning ( $A$  for advantage) and SARSA (State-Action-Response-State-Action).


## Reinforcement Learning: A Brief Introduction

Reinforcement learning is characterized by a sequence of interactions between a learning agent and the environment it wants to learn about. At every decision point and stage, the agent observes a certain states of the environment, and chooses an action (make a decision) from a set of possible actions. The environment responds to the actions by making a transition to a new state. In addition to observing the new state, the agent observes a reward that is meant to assess the immediate desirability of the action chosen by the agent. *State*, *action* and *reward* are the three basic elements of the reinforcement learning framework. Here are some illustrative examples of reinforcement learning (Sutton and Barto 1998, 6):

- A master chess player makes a move. The choice is informed both by planning – anticipating possible replies and counter-replies – and by immediate, intuitive judgments of the desirability of particular position and moves.
- An adaptive controller adjusts parameters of a petroleum refiner’s operation in real time. The controller optimizes the yield/cost/quality tradeoff on the basis of specified marginal costs without sticking strictly to the set points originally suggested by engineers.
- A gazelle calf struggles to its feet minutes after being born. Half an hour later it is running in 20 miles per hour.

The most widely-known context where the reinforcement learning is applied is Markov Decision Process. In a Markov Decision Process setting, the probability

Figure 6.1: Reinforcement Learning in Markov Decision Process



of the environment making a transition to a new state, given the current state and action, does not depend on the distant past of the environment (the Markov property of memorylessness), which Figure 6.1 illustrates. In this setting, the goal of reinforcement learning is to learn how to map states to actions so as to maximize the total expected future reward (Szepesvari 2010). The reward itself is usually a random variable, and hence the goal can be formulated in terms of an expectation (Sutton and Barto 1998, 61).

To use a game-theoretic terminology, the number of stages in a reinforcement learning problem can be either finite or infinite. And the future reward can either be with or without discounting ( $\delta$ ).<sup>3</sup> To simplify the illustration, let me consider only the finite-horizon without discount. In reinforcement learning, a *policy* defines the agent's behavior, i.e., which action to be taken based on the current state, at any given stage. A deterministic policy is a vector of mappings, with as many

---

<sup>3</sup>The discount factor,  $\delta$ , determines the importance of future rewards. A factor of 0 indicates that the agent only considers current short-term rewards (myopic), while a factor approaching 1 that the agent values the future reward almost same as the current one.

as the number of stages, where each component of the vector is a mapping from the state space (set of possible states) to the action space (set of possible actions) corresponding to a stage. In general, a policy can be stochastic. In stochastic policy, the mappings are from the stage-specific state spaces to the space of probability of distributions over the stage-specific action spaces. The discussion in this chapter is confined to deterministic policy.

Another key concept that is of crucial importance in policy evaluation is the *values function* or values (Sutton and Barto 1998, 68). While rewards reflect immediate desirability of an action, values represent the long run benefit. To summarize, the values of a given state, with respect to a given policy, is the total expected future reward of an agent, starting with that state, and following the given policy to select actions thereafter. Thus, the goal of reinforcement learning is to estimate a policy that maximizes the values over a specified class of policies. Since the value is a function of the state and policy, the maximization of values over policy space can happen for each state, or be averaged over all the states (Sutton and Barto 1998, 90).

## **Probabilistic Framework for Reinforcement Learning**

In the reinforcement learning framework, the agent and the environment interacts at each of a possible infinite number of stages in many sequential decision problems, and has the strong similarity to the dynamic programming (Parmigiani and Inoue 2009). However, again, for convenience of illustration, I will focus on the finite horizon ( $K$ ) problems without discounting. The setup and notations in this

section are largely similar to what I have already discussed in Chapter 5.

At stage  $j$  ( $1 \leq j \leq K$ ), the agent observes a state  $O_j \in O_j$  and takes an action  $A \in A_j$ , where  $O_j$  is the state space and  $A_j$  is the action space. To make the settings similar to those in causal inference, let us restrict the discussion to settings where  $O_j$  can be a vector consisting of discrete or continuous variables, but the action  $A_j$  can only be discrete. Partly as a consequence of its action, the agent receives a real-valued reward  $T_j \in \mathbb{R}$ , and moves on to the next stage with a new state  $O_{j+1} \in O_{j+1}$ . Using the same notations as in the previous chapters, define  $\bar{O}_j \equiv (O_1, \dots, O_j)$  and  $\bar{A}_j \equiv (A_1, \dots, A_j)$ . Also define the history  $H_j$  at stage  $j$  as the vector  $(\bar{O}_j, \bar{A}_{j-1})$ . At any stage  $j$ , the quantities  $O_j, A_j, Y_j$  and  $H_j$  are random variables, the observed values of which will be denoted by  $o_j, a_j, y_j$ , and  $h_j$ , respectively. The reward  $Y_j$  is conceptualized as a known function of the history  $H_j$ , the current action  $A_j$ , and the next state  $O_{j+1}$ . Thus,

$$Y_j = Y_j(H_j, A_j, O_{j+1}) = Y_j(\bar{O}_j, \bar{A}_j, O_{j+1}).$$

In some settings, there may be only one terminal reward  $Y_K$ ; rewards at all previous stages are taken to be 0. Here, rewards can easily be understood as outcomes in causal inference and the payoffs in decision theory.

Define a deterministic policy  $d \equiv (d_1, \dots, d_K)$  as vector of decision rules, where for  $1 \leq j \leq K$ ,  $d_j : H_j \rightarrow A_j$  is a mapping from the history space  $H_j$  to the action space  $A_j$ . As I mentioned above, a policy is stochastic, rather than deterministic, if the above mappings are from the history space  $H_k$  to the space of

probability distributions over the action space  $A_j$  which can be denoted  $d_j(a_j|h_j)$ . The collection of policies, depending on the history-space and action-space, defines a function space called policy space and is often denoted by  $D$  (Sutton and Barto 1998, 90)

The trajectory in the finite horizon, or the sequence of events (data) in causal inference, is composed of the set  $\{O_1, A_1, O_2, \dots, A_K, O_{K+1}\}$ . In the problem constructing dynamic treatment regimes, the data consist of the treatment records of  $n$  subjects, i.e.,  $n$  trajectories. It is assumed that the subjects are sampled at random according to a fixed distribution denoted by  $P_\pi$ . This distribution is composed of the unknown distribution of each  $O_j$  conditional on  $(H_{j-1}, A_{j-1})$ , and a fixed exploration policy for generating the action. Call the foregoing unknown conditional densities  $\{f_1, \dots, f_K\}$ , and denote the exploration policy by  $\pi = (\pi_1, \dots, \pi_K)$ , where the probability that action  $a_j$  is taken given history  $H_j$  is  $\pi_j(a_j|H_j)$ , which is the propensity score. Assume that  $\pi(a_j|h_j) > 0$  for each action  $a_j \in A_j$  and for each possible value  $h_j$ ; that is, all actions have a positive probability of being taken. Then, the likelihood under  $P_\pi$  of the trajectory  $\{o_1, a_1, o_2, \dots, a_K, o_{K+1}\}$  is

$$f_1(o_1)\pi_1(a_1|o_1) \prod_{j=2}^K f_j(o_j|h_{j-1}, a_{j-1})\pi_j(a_j|h_j)f_{K+1}(o_{K+1}|h_K, a_K)$$

Denote the expectation with respect to the distribution  $P_\pi$  by  $E_\pi$ . Let  $P_d$  denote the distribution of a trajectory where an arbitrary policy  $d = (d_1, \dots, d_k)$  is used to generate the actions. If  $d$  is a deterministic policy, then the likelihood under  $P_d$  of the trajectory  $\{o_1, a_2, o_2, \dots, a_K, o_{K+1}\}$  is



$$f_1(o_1)I[a_1 = d_1(o_1)] \prod_{j=1}^K f_j(o_j|h_{j-1}, a_{j-1})I[a_j = d_j(h_j)]f_{K+1}(O_{K+1}|h_K, a_K),$$

where  $I(\cdot)$  is an indicator function. For a stochastic policy  $d$ , the likelihood becomes

$$f_1(o_1)d_1(a_1|o_1) \prod_{j=1}^K f_j(o_j|h_{j-1}, a_{j-1})d_j(a_j|h_j)f_{K+1}(o_{K+1}|h_K, a_K)$$

Denote the expression with respect to the distribution  $P_d$  by  $E_d$ . The primary goal here is to estimate the optimal policy, say  $d^*$ , from the data on  $n$  finite horizon trajectories, which are not necessarily generated by the optimal policy. Optimal policy means the one with greatest possible value.

The value function for a state  $o_1$  with respect to arbitrary policy  $d$  is

$$V^d(o_1) = E_d\left[\sum_{j=1}^K Y_j(H_j, A_j, O_{j+1})|O_1 = o_1\right]$$

This represents the total expected future reward starting at a particular state  $o_1$  and thereafter choosing actions to the policy  $d$ . Given a policy  $d$ , the stage  $j$  value function for a history  $h_j$  is the total expected future rewards from stage  $j$  onward, and is given by

$$V_j^d(h_j) = E_d\left[\sum_{k=j}^K Y_k(H_k, A_k, O_{k-1})|H_j = h_j\right], \quad 1 \leq j \leq K$$

Note that, by definition,  $V_1^d(\cdot) = V^d(\cdot)$ . For convenience, set  $V_{K+1}^d(\cdot) = 0$ .

Then the values functions can be expressed recursively as follows:

$$V_j^d(h_j) = E_d\left[\sum_{k=j}^K Y_K(H_k, A_k, O_{k+1}) \mid H_j = h_j\right] \quad (6.1)$$

$$= E_d[Y_j(H_j, A_j, O_{j+1}) \mid H_j = h_j] + E_d\left[\sum_{k=j+1}^K Y_k(H_k, A_k, O_{k+1}) \mid H_j = h_j\right] \quad (6.2)$$

$$= E_d[Y_j(H_j, A_j, O_{j+1}) \mid H_j = h_j] \quad (6.3)$$

$$+ E_d\left[E_d\left[\sum_{k=j+1}^K Y_K(H_k, A_k, O_{k+1}) \mid H_{k+1}\right] \mid H_j = h_j\right] \quad (6.4)$$

$$= E_d[Y_j(H_j, A_j, O_{j+1}) \mid H_j = h_j] + E_d[V_{j+1}^d(H_{j+1}) \mid H_j = h_j] \quad (6.5)$$

$$= E_d[Y_j(H_j, A_j, O_{j+1}) + V_{j+1}^d(H_{j+1}) \mid H_j = h_j], \quad 1 \leq j \leq K \quad (6.6)$$

The optimal stage  $j$  value function for a history  $h_j$  is now defined as:

$$V_j^{opt}(h_j) = \max_{d \in D} V_j^d(h_j)$$

The optimal value function satisfies the Bellman equation.

$$V_j^{opt}(h_j) = \max_{a_j \in A_j} E[Y_j(H_j, A_j, O_{j+1}) + V_{j+1}^{opt}(H_{j+1}) \mid H_j = h_j, A_j = a_j],$$

when all observations and actions are discrete.<sup>4</sup>

---

<sup>4</sup>The Bellman equation also holds for more general scenarios, but with additional assumptions.

Finally, the marginal value of policy  $d$ , written  $V^d$ , is the average value function under that policy, averaged over possible observations, i.e.,

$$V^d = E_{O_1}[V^d(O_1)] = E_d\left[\sum_{k=1}^K Y_k(H_k, A_k, O_{k+1})\right]$$

Note that the above expectations is taken with respect to entire likelihood of the data, for the case of deterministic or stochastic policy, respectively. Thus the value of a policy is simply the marginal mean outcomes under that policy.

The primary statistical goal for a policy is to estimate its value, and a related problem would be to compare the values of multiple policies. In other words, the process is same as comparing mean outcomes of two or more (static) treatments in causal inference.

In many classical reinforcement learning, researchers often seek to estimate a policy that maximizes the value, i.e., optimal policy. One approach is to first specify a policy space, and then employ some methods to estimate the value of each policy in that space to find the best one. An alternative approach is to work with what is known as an action-value function, or simply a  $Q$ -function instead of the value function  $V^d$  defined above.  $Q$ -functions are defined as follows:

$$Q_j^d(h_j, a_j) = E[Y_j(H_j, A_j, O_{j+1}) + V_{j+1}^d(H_{j+1}) | H_j = h_j, A_j = a_j]$$

The optimal stage  $j$   $Q$ -function is

$$Q_j^{opt}(h_j, a_j) = E[Y_j(H_j, A_j, O_{j+1}) + V_{j+1}^{opt}(H_{j+1}) | H_j = h_j, A_j = a_j]$$

Now the major task lies in estimating  $Q_j^{opt}$ , since this can directly lead to the optimal policy.

## Estimating Optimal Dynamic Treatment Regime

The typical approach to estimate the optimal policy, known as policy search method in reinforcement learning, is to first specify a policy space,  $D$ , and then adopt any suitable method to estimate the value of each candidate policy  $d \in D$  to find the best one, say  $\hat{d}^{opt}$ . More precisely,

$$\hat{d}^{opt} = \arg \max_{d \in D} \hat{V}^d$$

As I mentioned in the previous chapter, there are two fundamentally different ways of handling dynamic treatment and time-varying covariates within the potential outcomes framework. One is the marginal procedure, where the aim is to estimate a treatment effect, and all other effects are considered nuisances to be adjusted for. The policy search method, which include inverse probability weighting, falls into this way of handling the dynamic treatment. While the policy search approach is typically non-parametric or semi-parametric, requiring only mild assumptions about the data, the main issue is the high variability of the value function estimates, and the resulting high variability in the estimated optimal policies (Chakraborty and Moodie 2013).

The other is joint modeling, where all or most components in the process are modeled. For example, the change of treatment or policy is explicitly modeled as a function of previous observations, as are also time-varying covariates.  $G$ -estimation, likelihood-based estimation, and the Bayesian approach can be included into this category.

For simplicity, let me first describe  $Q$ -learning for two stages (intervals), and then generalize it to  $K(\geq 2)$  stages. In a two-stage case, the data on a single subject are given by the sequence or trajectory  $(O_1, A_1, O_2, A_2, O_3)$ . The histories are given by  $H_1 \equiv O_1$  and  $H_2 \equiv (O_1, A_1, O_2)$ . The order of the data are same as in dynamic treatment regime in Chapter 5. The data available for estimation consist of a random sample of  $n$  subjects. For simplicity, assume the data are drawn from sequentially randomized trials with two possible treatments in each stage,  $A_j \in \{0, 1\}$  and that they are randomized (conditional on history) with known randomization probabilities. The study can have either a single terminal reward (primary outcome),  $Y$ , observed at the end of stage 2, or two rewards (intermediate and final outcomes),  $Y_1$  and  $Y_2$ , observed at the end of each stage. The case of single terminal outcome  $Y$  is viewed as a special case with  $Y_i \equiv 0$  and  $Y_2 \equiv Y$ . A two-stage policy consists of two decision rules, say  $(d_1, d_2)$ , with  $d_j(H_j) \in \{0, 1\}$ .

One simple method to construct optimal dynamic treatment regime  $d^{opt} \in (d_1^{opt}, d_2^{opt})$  is  $Q$ -learning. First define the optimal  $Q$ -function for the two stages as follows:

$$Q_2^{opt}(H_2, A_2) = E[Y_2|H_2, A_2], \quad (6.7)$$

$$Q_1^{opt}(H_1, A_1) = E[(Y_1) + \max_{a_2} Q_2^{opt}(H_2, a_2)|H_1, A_1] \quad (6.8)$$

If the above two  $Q$  functions were known, the optimal dynamic treatment regime  $(d_1^{opt}, d_2^{opt})$ , using a backwards induction, would be

$$d_j^{opt}(h_j) = \arg \max_{a_j} Q_j^{opt}(h_j, a_j), \quad j = 1, 2$$

The true  $Q$ -functions are not known and need to be estimated from the data. Note that  $Q$ -functions are conditional expectations, and hence a natural way of modeling them is by regression models (Buşoniu et al. 2012). Consider linear regression models for the  $Q$ -function. Let the stage  $j$  ( $j = 1, 2$ )  $Q$  function be modeled as

$$Q_j^{opt}(H_j, A_j; \beta_j, \psi_j) = \beta_j^T H_{j0} + (\psi_j^T H_{j1}) A_j,$$

where  $H_{j0}$  and  $H_{j1}$  are two (possibly different) vector summaries of the history  $H_j$ , with  $H_{j0}$  denoting the main effect of history, since the vector  $H_j$  also includes a term that corresponds to the main effect of treatment, similar to the intercept in regressions. To use the terms in epidemiology, the variables  $H_{j0}$  are often called predictive, while  $H_{j1}$  contains prescriptive or tailoring variables, which are triggering the changes of treatments.<sup>5</sup> To summarize, the  $Q$ -learning algorithm typically involves

---

<sup>5</sup>Triggering variables are the ones used to trigger a change in the treatment. In clinical trials,

the following two stage least square regressions (Sutton and Barto 1998; Buşoni et al. 2012; Nahum-Shani et al. 2012; Chakraborty and Moodie 2013):

1. Stage 2 regression:  $(\hat{\beta}_2, \hat{\psi}_2) = \arg \min_{\beta_2, \psi_2} \frac{1}{n} \sum_{i=1}^n (Y_{2i} - Q_2^{opt}(H_{2i}, A_{2i}; \beta_2, \psi_2))^2$
2. Stage 2 pseudo-outcome:  $\hat{Y}_{1i} = Y_{1i} + \max_{a_2} Q_2^{opt}(H_{2i}, a_2; \hat{\beta}_2, \hat{\psi}_2), i = 1, \dots, n.$
3. Stage 1 regression:  $(\hat{\beta}_1, \hat{\psi}_1) = \arg \min_{\beta_1, \psi_1} \frac{1}{n} (\hat{Y}_{1i} - Q_1^{opt}(H_{1i}, A_{1i}; \beta_1, \psi_1))^2$

In the step 2, the quantity  $\hat{Y}_{1i}$  is a predictor of the unobserved random variable  $Y_{1i} + \max_{a_2} Q_2^{opt}(H_{2i}, a_2), i = 1, \dots, n.$  Once the  $Q$ -function have been estimated, finding the optimal dynamic treatment regimes is simple. The estimated optimal dynamic treatment regime using  $Q$ -learning is given by  $(d_1^{opt}, d_2^{opt})$ , where the stage  $j$  optimal rule is specified as

$$\hat{d}_j^{opt}(h_j) = \arg \max_{a_j} Q_j^{opt}(h_j, a_j; \hat{\beta}_j, \hat{\psi}_j), j = 1, 2$$

This procedure can easily be generalized to  $K > 2$  stages. Define  $Q_{K+1}^{opt} \equiv Q$ , and

$$Q_j^{opt}(H_j, A_j) = E(Y_j + \max_{a_{j+1}}^{opt}(H_{j+1} + a_{j+1}) | H_j, A_j], \quad j = 1, \dots, K$$

Stage specific  $Q$ -function can be parametrized as before,

---

these can be things like early signs of non-response, manifestation of side effects, or environmental or social characteristics. The idea is to identify the variables that best indicate when the appropriate treatment has changed.

$$Q_j^{opt}(h_j, A_j; \beta_j, \psi_j) = \beta_j^T H_{j0} + (\psi_j^T H_{j1}) A_j, j = 1, \dots, K$$

Finally, for  $j = K, K - 1, \dots, 1$ , moving backward through stages, the regression parameters can be estimated as

$$(\hat{\beta}_j, \hat{\psi}_j) = \arg \min_{\beta_j, \psi_j} \frac{1}{n} \sum_{i=1}^n (Y_{ji} + \max_{a_{j+1}} Q_{j+1}^{opt}(H_{j+1}, a_{j+1}; \hat{\beta}_{j+1}, \hat{\psi}_{j+1}) - Q_j^{opt}(H_{ji}, A_{ji}; \beta_j, \psi_j))^2,$$

where  $(Y_{ji} + \max_{a_{j+1}} Q_{j+1}^{opt}(H_{j+1}, a_{j+1}; \hat{\beta}_{j+1}, \hat{\psi}_{j+1}))$  is the stage  $j$  pseudo outcome. As before, the estimated optimal dynamic treatment regime is given by  $(\hat{d}_1^{opt}, \dots, \hat{d}_K^{opt})$ , where

$$\hat{d}_j^{opt}(h_j) = \arg \max_{a_j} Q_j^{opt}(h_j, a_j; \hat{\beta}_j, \hat{\psi}_j), j = 1, \dots, K$$

Alternatively, the following one-step procedure for estimating the optimal dynamic treatment regime could be more natural. In this approach, one would model the conditional mean outcome  $E(Y|O_1, A_1, O_2, A_2)$  and run regression, where the estimated optimal policy would be

$$(\hat{d}_1^{opt}, \hat{d}_2^{opt}) = \arg \max_{a_1, a_2} E(Y|o_1, a_1, o_2, a_2).$$

However, the problem in this formulation is the possible bias in the estimation of stage 1 treatment effect, which arises as a consequence of what is known as collider-stratification bias or Berkson's paradox, as I discussed in Chapter 4.



## Simulation

In the previous chapters, I have used the propensity score and the inverse probability weight as devices to adjust for time-varying confounding and censoring in Chapter 5, using the transitional justice’s effects as examples. In this section, I run a series of simulation to examine the relative efficiency of the estimators that we are most interested in: inverse probability weight. The primary focus of this simulation is to examine how inverse probability weight estimator performs compared to other more commonly used estimators.

In this simulation, I distinguish between the components  $O_j$  that are tailoring variables, which trigger the change of the treatments, and the predictive variables including potential confounders, denoted  $X_j, j = 1, 2$ . Three of the adjustment methods adapt  $Q$ -learning by re-defining the history vectors,  $H_1$  and  $H_2$ . The fourth approach uses a propensity score matching, while the fifth relies on the inverse probability weighting. If we denote the interval-specific propensity score ( $\pi$ ) by  $\pi_1 = \Pr(A_1 = 1|X_1), \pi_2 = \Pr(A_2 = 1|X_1, X_2)$ , then we can compare  $Q$ -learning using the following five implementations:

1. Include only  $O_j$ :  $H_1 = O_1; H_2 = O_2$ .
2. Include all covariates as linear terms:  $H_1 = (X_1, O_1), H_2 = (X_1, X_2, O_1, A_1, O_2)$ .
3. Include the propensity scores for each stage as linear terms:  $H_1 = (\pi_1, O_1), H_2 = (\pi_2, O_1, A_1, O_2)$ .
4. Use propensity score matching with  $H_1 = O_1; H_2 = O_2$ .

5. Use inverse probability treatment weighting with  $H_1 = O_1; H_2 = O_2$ .

Next, I construct multiple simple scenarios where the time-varying confounding is problematic. Figure 6.2 shows the DAG for the following models.

1. Model A: A single continuous confounder,  $X_j$ , exists at each interval, where  $X_1 \sim N(0, 1)$  and  $X_2 \sim N(\eta_0 + \eta_1 X_1, 1)$ . And set  $\eta_0 = -0.5, \eta_1 = 0.5$ . Treatment assignment depends on the values of confounding variable:  $\Pr(A_j = 1|X_j) = 1 - \Pr(A_j = 0|X_j) = \text{expit}(\zeta_0 + \zeta_1 X_j), j = 1, 2$ . The binary covariates that interact with treatment to produce a subject-specific rule are generated by equal probability (1/2).

$$\Pr(O_2 = 1|O_1, A_1) = 1 - \Pr(O_2 = 0|O_1, A_1) = \text{expit}(\delta_1 O_1 + \delta_2 A_1).$$

Then, we get a regression model:


$$Y = \gamma_0 + \gamma_1 X_1 + \gamma_2 O_1 + \gamma_3 A_1 + \gamma_4 O_1 A_1 + \gamma_5 X_2 + \gamma_6 A_2 + \gamma_7 O_2 A_2 + \gamma_8 A_1 A_2 + \varepsilon.$$

Here, parameters were chosen to produce regular setting,<sup>6</sup> where asymptotic distributions of estimator converge uniformly:  $\gamma = (0, \gamma_1, 0, -0.5, 0, \gamma_5, 0.25, 0.5, 0.5)$  and  $\delta = (0.1, 0.1)$ . I begin with the model where  $\zeta_0 = \zeta_1 = 0$  and  $\gamma_1 = \gamma_5 = 0$ . This is a randomized trial setting.

---

<sup>6</sup>Non-regular settings mean that asymptotic distributions of estimator of the treatment effect parameter does not converge uniformly over the parameter space (Chakraborty and Moodie 2013, 128).

Figure 6.2: Directed Acyclic Graph for Simulated Data



2. Model *B*: Now set the values of  $\gamma_1$  and  $\gamma_5$  at 1. This is also a randomized trial setting.
3. Model *C*: Based on Model *B*, now set  $\zeta_0 = 0.8$ ,  $\zeta_1 = 1.25$ . In this setting, treatment is now confounded by  $X_1$  and  $X_2$ .
4. Model *D*: Confounders  $X_1$  and  $X_2$  are now binary, instead of continuous. Set  $\Pr(X_1 = 1) = 1 - \Pr(C_1 = 0) = 1/3$  and  $\Pr(X_2 = 1) = 1 - \Pr(X_2 = 0) = \text{expit}(0.6 \cdot X_1)$ .  $\gamma_1 = \gamma_5 = 1$ .
5. Model *E*: This setting is same as Model *D*, except that  $\gamma_1 = 0$  so that  $X_2$  is the predictor of  $Y$ , but  $X_1$  is not.

As in the  $g$ -estimation, the focus here is the parameter  $\psi_{10}$ , the parameter for the first-stage  $Q$ -function which corresponds to the main effect of  $A_1$ . Performance of the six different  $Q$ -learning approaches are given Table 6.1 for sample size 250. Note that I did not examine propensity score matching here when the confounders are binary, although they are not impossible as I used for multiple treatments in

Chapters 4 and 5 (Imbens 2000; Imai and van Dyk 2004; Egger and von Ehrlich 2013).

Table 6.1 shows that all methods perform well with random allocation of treatment (Models *A* and *B*), with better performance when  $X_j$  does not predict the outcome. Obviously, correct specification is not crucial in the randomized trial setting.

In settings where treatment is confounded by the continuous confounders  $X_1$  and  $X_2$  (Model *C*), only covariate adjustment provides unbiased estimates. The same pattern holds when confounders are binary (Model *D*), with one exception: if there exists a single confounder at each interval and only  $X_2$ , but, not  $X_1$ , affects  $Y$  then including  $\pi$  in the  $Q$ -function model performs as well as including  $X_2$  in the model, since propensity scores act as a re-scaled version of  $X_2$ . Although regression-based methods yield unbiased estimates of the parameters associated with treatment  $A_2$ , i.e., the variables contained in  $H_{21}$ , the methods do not yield a good prediction of the stage 1 pseudo outcome itself, since the model mis-specifies the functional form of the dependence of that pseudo-outcome on important predictors  $X_1$  and  $X_2$ . This leads to bias in the stage 1 parameter estimates. Further, note that the matching estimator targets the average treatment effect on the treated (ATT), and has increased variability because it matched the pairs with replacement.

In the simulation above, the data were generated in such a way that a model that includes the confounding variables as linear terms in the  $Q$ -function was correctly specified.

Table 6.1: Performance of Four Adjustment Methods

	Bias	MSE	Variance
Model A: $\zeta_0 = \zeta_1 = \gamma_1 = \gamma_5 = 0$			
null	-0.000	0.009	0.009
linear model	-0.000	0.009	0.009
$\pi$ as linear terms	-0.000	0.009	0.009
$\pi$ matching	-0.009	0.017	0.017
Inverse probability Weigh	-0.001	0.009	0.009
Model B: $\gamma_1 = \gamma_2 = 1$			
null	0.008	0.032	0.032
linear model	0.009	0.011	0.011
$\pi$ as linear terms	0.006	0.019	0.019
$\pi$ matching	-0.015	0.047	0.047
Inverse Probability Weight	0.008	0.023	0.023
Model C: $\zeta_0 = 0.8, \zeta_1 = 1.25$			
null	-0.710	0.534	0.025
linear model	-0.003	0.011	0.012
$\pi$ as linear term	-0.263	0.089	0.024
$\pi$ matching	-0.265	0.143	0.068
Inverse Probability Weight	-0.440	0.214	0.020
Model D: binary, $\gamma_1 = \gamma_5 = 1$			
null	-0.607	0.388	0.023
linear model	0.010	0.014	0.014
$\pi$ as linear terms	-0.271	0.098	0.029
inverse probability weight	-0.337	0.132	0.025
Model E: $\gamma_1 = 0, \gamma_1 = 1$			
null	-0.238	0.070	0.010
linear model	0.008	0.011	0.011
$\pi$ as linear terms	0.006	0.013	0.013
inverse probability weight	-0.137	0.030	0.011

The setup in the next simulation is for the scenario that will allow us to examine the performance of the adjustment methods without the additional complication of incorrect model specification. This is possible by generating data in which confounding is introduced by making the treatment assignment is based on the potential outcomes, i.e., which are not observed. As I mentioned in Chapter 2, in observational

studies, this is possible only in crossover design in which the effects from previous interval do not carry over. The data are created by generating the outcome under each of the four potential outcome sequences,  $(0, 0)$ ,  $(0, 1)$ ,  $(1, 0)$  and  $(1, 1)$ . The data generating steps are as follows:

1. Generate the first-stage variable  $O_1$ , using the binomial distribution,  $\Pr(O_1 = 1) = \Pr(O_1 = 0) = \frac{1}{2}$ .
2. Generate the potential values of the second-stage variable,  $O_2(A_1)$ , using  $\Pr(O_2 = 1|O_1, A_1) = 1 - \Pr(O_2 = 0|O_1, A_1) = \text{expit}(\delta_1 O_1 + \delta_2 A_1)$  for each possible value of  $A_1$ , which generates the potential second-stage value that would occur under each of  $A_1 = 0$  and  $A_1 = 1$ .
3. Generate the vector of potential outcome,  $\mathbf{Y} = \gamma_0^* + \gamma_1^* \mathbf{O}_1 + \gamma_2^* \mathbf{A}_1 + \gamma_3^* \mathbf{O}_1 \mathbf{A}_1 + \gamma_4^* \mathbf{A}_2 + \gamma_5^* \mathbf{O}_2 \mathbf{A}_2 + \gamma_6^* \mathbf{A}_1 \mathbf{A}_2 + \varepsilon$ , where  $\varepsilon$  is a multivariate normal error term with mean  $(0, 0, 0, 0)^\top$  and a covariance matrix that takes the value of 1 on its diagonal and 0.5 on all off-diagonals. In the above equation,  $\mathbf{O}_1$  is the set of  $4 \times 1$  vectors consisting of  $O_1$  from the first step repeated four times  $\mathbf{A}_1 = (0, 0, 1, 1)$ ,  $\mathbf{O}_2 = (O_2(0), O_2(0), O_2(1), O_2(1))$  using the potential values generated in the second step and  $\mathbf{A}_2 = (0, 1, 0, 1)$ .
4. Set the confounders to be  $X_1 = \bar{\mathbf{Y}}$  and  $X_2 = \max(\mathbf{Y})$ .
5. From the four possible treatment paths and corresponding potential outcomes, select the observed data using  $\Pr(A_j = 1|X_j) = 1 - \Pr(A_j = 0|X_j) = \text{expit}(\zeta_0 + \zeta_1 X_j)$ ,  $j = 1, 2$ .

The vector of  $\delta$ s was set to  $(0, 1, 0, 1)$ , while the vector of  $\gamma^*$ s was taken to be  $(0,0,-0.5,0,0,0.25,0.5,0.5)$ . In simulations where treatment is randomly allocated,  $\zeta_0 = \zeta_1 = 0$ , while for confounded treatment,  $\zeta_0 = 0.2, \zeta_1 = 1$ . Since the  $Q$ -functions will not depend on the values of  $X_1$  and  $X_2$ , any model for  $Q$ -function that includes  $O_1, A_1, O_2, A_2$  and the appropriate interactions are considered correctly specified. However, the observed treatment depends on  $X_1$  and  $X_2$ , which are functions of the potential outcomes.

Table 6.2 shows the simulation results. We can see that all methods of adjusting for confounding provide considerably improved estimates in terms of bias for small samples, but in large samples ( $n=1,000$ ), only inverse probability weighting or directly adjusting for covariates by including them as linear terms in the model for the  $Q$ -function provide the required removal of bias. While these simulations provide a useful demonstration of the methods of adjustment in principle, these results cannot easily be generalized to real data because it is difficult to conceive of a situation in which counterfactual outcomes could be measured and used as covariates.

Table 6.2: Performance of Adjustment Methods under Confounding

	Randomized				Confounded			
	Bias	MSE	Variance	Coverage	Bias	MSE	Variance	Coverage
<i>n</i> = 500								
null	0.002	0.007	0.007	94.1	0.229	0.061	0.008	26.5
linear model	0.001	0.002	0.002	95.2	0.004	0.004	0.004	92.7
$\pi$ linear term	0.001	0.005	0.004	96.0	0.048	0.009	0.006	89.0
$\pi$ matching	0.002	0.010	0.011	98.0	0.138	0.030	0.010	75.7
IPW	0.000	0.004	0.004	94.0	0.011	0.008	0.008	92.0
<i>n</i> = 1,000								
null	-0.001	0.002	0.002	93.4	0.225	0.053	0.002	0.8
linear model	0.000	0.010	0.001	93.5	0.004	0.001	0.001	92.5
$\pi$ linear term	-0.000	0.001	0.001	95.5	0.044	0.004	0.002	78.0
$\pi$ matching	-0.002	0.003	0.003	97.5	0.136	0.018	0.003	32.0
IPW	-0.001	0.001	0.001	93.6	0.002	0.002	0.002	93.3

MSE: Mean Squared Error, Coverage: Coverage of 95% confidence interval

## Application: Economy and Human Rights Trials

In this section, I construct a simple  $Q$ -learning model with transitional justice, physical integrity and GDP per capita. The main question is to find the optimal dynamic treatment regime in the presence of time-varying confounding of the covariates. In this estimation, the following variables are used:  $A_1$ : Human rights trials in the first three years after transition,  $A_2$ : Human rights trials in the year 3-5 after democratic transition,  $H_1$ : GDP per capita (logged). The outcome variable is physical integrity rights. Table 6.3 shows the regression coefficients for the two stage regression for  $Q$ -learning, and Table 6.4 the point estimates and the confidence intervals for contrasts for the policies in the  $Q$ -learning.<sup>7</sup>

<sup>7</sup>An R package, `qLearn`, written by Jingyi Xin, Bihbas Chakraborty, and Eric B. Laber, was used for part of the estimation.



Table 6.3: Coefficients in Two Stage Regression for  $Q$ -learning

Variable	Coef. (stage 2)	Coef. (stage 1)
$H_1$	4.775	-0.261
$A_1$	-2.080	-5.109
$A_2$	-3.105	
$H_1 \times A_1$	0.166	0.636
$A_1 \times A_2$	0.940	
$A_2 \times H_2$	0.411	

Table 6.4: Estimates and Confidence Intervals for Contrasts in Two Stages

Stage	Contrast	Estimates	Lower Bound	Upper Bound
1	1	4.78	1.64	7.07
	2	-0.20	-0.52	0.23
	3	-2.08	-4.76	0.65
	4	-3.10	-6.19	0.33
	5	0.17	-0.21	0.59
	6	0.94	-0.42	2.69
	7	0.41	-0.02	0.79
2	1	8.85	0.25	15.25
	2	-0.26	-1.26	0.75
	3	-5.11	-12.56	3.54
	4	0.64	-0.41	1.65

## Discussion

I have shown that the dynamic treatment regimes can be seamlessly combined with reinforcement learning, which is relatively less known outside of computer science. Using the basic terminology of reinforcement learning, especially  $Q$ -learning, I illustrated the similarity between the causal inference and reinforcement learning, by simply converting the key terms such as treatment regime, covariates, and treatment in causal inference into policy, state space and action space in reinforcement learning, among others.

Continuing my discussion of inverse probability weights in the previous chapters, I showed that the inclusion of inverse probability weights and the linear modeling could be far more effective in  $Q$ -learning adjustment than other propensity score based methods, including matching, with simulated data, and that the optimal policy under  $Q$ -learning can simply be implemented by running a series of regressions recursively. In the last section, I also demonstrated the  $Q$ -learning with linear regression can be applied in finding the optimal policy for human rights trials.

The implication of my discussion in this chapter is not clear, because the question is whether these ideas can be generalized to more complicated scenarios, let alone to the real observational data in political science, because the setup of the simulation is based on extremely simple situations. To re-emphasize, the discussion of  $Q$ -learning and the comparison of a series of estimation methods with simulated data show the difficulty of bias reduction in observational data, and the importance of specification. For that purpose, we need to have the data for *all* relevant covariates at every interval, and *all* intermediate decisions so that there is no unmeasured

confounding. Unlike in randomized trials, those data are not usually available in social science, unless the researcher collects the data with a clear intent from the beginning. The data requirement also limits its applicability in comparative politics, including the studies on transitional justice. Ultimately, we need to have the correct functional form and knowledge that could link the covariates and treatments.

## Chapter 7

### Conclusion

In this dissertation, I addressed some of the methodological problems in causal inference within the potential outcomes framework. Those problems pertain to whether and how we can estimate the causal effects when the assumptions for identifications are violated and/or when the various sorts of confounding, including time-varying one, exist. To deal with those problems, I introduced statistical and conceptual tools that are not well-known in political science, which include dynamic treatment regime, inverse probability weight,  $g$ -estimation and reinforcement learning,  $Q$ -learning, among others. These tools originate in epidemiology, biostatistics or computer science to deal with particular problems in each discipline, but could also be usefully incorporated into the toolkits of political scientists to address causality with observational data.

The idea of dynamic treatment regime, which originates in clinical trials and epidemiology to deal with the chronic problems and to provide personalized medicine, offers useful conceptual tools to estimate the long-term effects on long-term problems. The existing literature on causal inference has mostly focused on the effects of a single shot treatment in static settings. The best strategy to deal with long-term problem in the existing framework is to select limited time frame and estimate the

causal effect in that arbitrarily selected frame. Instead, dynamic treatment regime suggests that we need to examine the sequence of treatment as a regime, in which the outcome at the last stage is the criteria for the optimal treatment. The idea of dynamic treatment regime provides a way to understand sequential treatments and decisions.

Most of social problems like inequality or human rights violations need long-term treatment, and the static framework of causal inference is not adequate to deal with these problems because the humans change their strategies and actions in response to the changing situations. There have been long-standing tradition to incorporate this aspect of human action in econometrics on selection, which could be combined with the idea dynamic treatment regime. This idea is adequate to modeling transitional justice, whose emphasis is moving from perpetrator- or victim-centered approach to a more holistic approach.

Propensity score and inverse probability weights could be very useful tools in observational studies in which randomization is not possible. These are especially useful to political scientists, especially in the field of comparative politics, who usually have very limited number of cases and, for that reason, the sophisticated techniques of balancing covariates such as matching are not be successful at the country-level analysis. In this situation, propensity scores and inverse probability weighting could be very useful, especially because they are based on the statistical techniques that are in the standard toolkits in political science such as OLS, GLM or GAM. I extended the inverse probability weights to the case of multi-level treatments, instead of binary one. These could be useful in estimating the effects when the treatments are multiple-

level. Many applied works of causal inference in political science arbitrarily confine their application to the case of binary treatment or arbitrarily making treatment as binary by combining different treatments. These are far from reality.

Structurally Nested Models and an accompanying  $g$ -estimation strategy, originally proposed by Robins, try to find the optimal sequential decision making by direct modeling. Although my application of  $g$ -estimation is limited to the simple scenarios with limited number of variables, the idea could be expanded to more complex situations with a longer time frame, to time-event data or the effects of timing on the outcome.

The decision-theoretic view on the causal inference makes it possible to understand causal inference without the potential outcomes framework. However, the biggest advantage of incorporating decision theoretic view into causal inference lies in the possibility of modeling the complex human action and estimating the interaction between the humans and their environments. In the typical causal inference literature, the human-being are not necessarily modeled as a smart player in the sense that he or she adjusts his or her behavior in response to the changing situations. To model this continuing interaction, the ideas of sequential decision making and dynamic programming, which deal with the problem of finding the optimal solution computationally, provide rich soil to model the human action.

I attempted to combine the ideas of dynamic treatment regime and reinforcement learning, and discussed a particular kind of reinforcement learning,  $Q$ -learning, as an example, using simulated data. As I mentioned in Chapter 6, inverse probability weights and linear regressions can easily be mobilized in estimating the causal

effects and finding the optimal treatment regimes.

Although the main focus of the dissertation is to describe and illustrate the methods to deal with time-varying confounding, I tried to compare the estimation of causal effects of transitional justice measures with or without adjustment of the confounding I discussed. Overall, they showed the stark difference in estimates and the stronger causal effects of more transitional justice in improving human rights and maintaining democracy.

## Implications

The typical causal question in social science takes the form whether  $x$  caused  $y$ , and does not take into the interactions between the intermediate treatments and covariates over time seriously, except the occasional use of the Bayesian updating (Bates, de Figueiredo, Jr. and Weingast 1998). The decision-maker monitors the intermediate outcomes and adjusts the strategies to maximize the long terms outcomes. The intuition is very common in everyday decision making, and game theory have developed some conceptual tools to deal with it. However, its implication with reference to quantitative methodology is not well understood, let alone developed to deal with the problems. Without dealing with this problem, we may still say  $x$  caused  $y$ , but the quantity of the causal effects could be biased due to the confounding. However, the emphasis must be placed on the relationship between  $x$  and  $y$ . Whether it is truly causal is not in the purview of causal inference literature.

The ideas of optimal dynamic treatment regime and reinforcement learning can be usefully combined with dynamic game theory and/or differential game (Haurie

and Zaccour 2005; Long 2010; Haunschmied, Veliov and Wrzaczek 2014; Isaacs 1965; Dockner et al. 2001). Counterfactual causal inference is the quantitative extension of the well-known most similar systems design in comparative politics, and propensity scores, the inverse probability weight and  $g$ -estimation provides the supplemental method in situations when confounding needs to be adjusted for, but other methods of balancing is not feasible.

The ideal way to deal with time-varying confounding is to model the confounding directly. However, in many circumstances the state of knowledge does not allow it, and the estimation of causal effects through weighting could be useful tools for adjusting selection bias.

The most serious practical problem in applying the idea of dynamic treatment regime is the paucity of cases and the lack of relevant data in political science, especially in comparative politics. First, even with a binary treatment, we have  $2^4 = 16$  combinations of treatments for four stages. Given that the treatment options are more diverse, the number of sequences and their combinations quickly increases. Second, the typical data in political science do not take the form relevant for analyzing the decision process, and are not useful for modeling for a dynamic treatment regime. In order to estimate the treatment effect in dynamic treatment regime, we need to have data on all covariates, intermediate and final treatment decisions, and intermediate and final outcomes so that there is no unmeasured confounding. Especially problematic in political science is the lack of data on the intermediate decisions and intermediate outcomes. For a better analysis of dynamic treatment regime, the researcher needs to collect the data with clear purpose.



## Future Research

Throughout the dissertation, I have delved into the questions of time-varying confounding in causal inference. During the process, I have mostly focused on relatively straightforward forms of treatment, binary or ordinal. However, the treatment can be expanded into other types of variables, e.g., continuous or count. We can, for example, examine the effect of time to event, say the interval between democratic transition and the first trial, on the human rights and democracy for the country. The similar idea can be expanded into examining the effect of timing, for example election, on the regime survival, using this framework. These will be the direct attempts to deal with sequence, timing, and path dependence using quantitative methods.

I relied on the parametric models in creating inverse probability weights such as OLS, logit, or multinomial logit. However, other non-parametric models or semi-parametric models such as GAM could easily be used in constructing weights when the functional form between the covariates and the treatment is not clear. However, the primary task is to understand the selection or censoring mechanism that could justify the functional forms in use.

I focused on the causal *effects* of some treatments, transitional justice, in this dissertation. By definition, treatment must be manipulable (no manipulation, no causation). However, there are many variables that are not manipulable, hence no treatment, but have causal power such as gender, race, or culture. The recent efforts to conceptualize causal attributes, instead of causal effects, are promising venues of research in expanding the applicability of potential outcome framework and, furthermore, in expanding the horizons of our understanding of causality.

According to Pearl, causality is the issue that many statisticians avoided because of its ambiguity and difficulty (Pearl 2009). It is evident that the potential outcomes framework could be useful to conceptualize the causality and provide many statistical tools such as matching of various sorts for that purpose. The study on causality has a long history in philosophy and economics, and there has been few attempt to link the causality in the economics such as granger-causality in time series with the potential outcomes framework or graphical causal model. This dissertation is a step toward combining and re-interpreting the old ideas with the new tools.

## Bibliography

- Aalen, Odd O., Ørnulf Borgan and Håkon K. Gjessing. 2008. *Survival and Event History Analysis: A Process Point of View*. New York: Springer.
- Achen, Christopher. 2001. “Why Lagged Dependent Variables Can Suppress the Explanatory Power of Other Independent Variables.” Unpublished.
- Anan, Kofi. 2004. The Rule of Law and Transitional Justice in Conflict and Post-conflict Societies. Technical report United Nations Security Council.
- Arenhövel, Mark. 2008. “Democratization and Transitional Justice.” *Democratization* 15(3):570–587.
- Backer, David. 2009. Cross-National Comparative Analysis. In *Assessing the Impact of Transitional Justice: Challenges for Empirical Research*, ed. Hugo van der Merwe, Victoria Baxter and Audrey R. Chapman. Washington, D.C.: United States Institute of Peace Press pp. 23–89.
- Barkan, Elazar. 2000. *The Guilt of Nations: Restitution and Negotiating Historical Injustices*. New York: Norton.
- Bates, Robert H., Avner Greif, Margaret Levi, Jean-Laurent Rosenthal and Barry Weingast. 1998. *Analytic Narratives*. Princeton: Princeton University Press.

- Bates, Robert H., Rui J.P. de Figueiredo, Jr. and Barry R. Weingast. 1998. "The Politics of Interpretation: Rationality, Culture, and Transition." *Politics and Society* 26(4):603–642.
- Beebe, Helen, Christopher Hitchcock and Peter Menzies, eds. 2012. *The Oxford Handbook of Causation*. New York: Oxford University Press.
- Bellman, Richard. 1957. *Dynamic Programming*. Princeton, N.J.: Princeton University Press.
- Berkson, Joseph. 1946. "Limitations of the Application of Fourfold Table Analysis to Hospital Data." *Biometrics Bulletin* 2(3):47–53.
- Berry, Scott M., Bradley P. Carlin, J. Jack Lee and Peter Müller. 2010. *Bayesian Adaptive Methods for Clinical Trials*. Boca Raton, FL: CRC Press.
- Bertsekas, Dimitri P. 2007. *Dynamic Programming and Optimal Control*. 3rd ed. Belmont, MA: Athena Scientific.
- Bingham, Richard D. and Claire L. Felbinger. 2002. *Evaluation in Practice: A Methodological Approach*. 2nd ed. Chatam House Publishers.
- Blackwell, Matthew. 2012. "A Framework for Dynamic Causal Inference in Political Science." *American Journal of Political Science* 57(2):504–520.
- Brady, Henry E. 2008. Causation and Explanation in Social Science. In *The Oxford Handbook of Political Methodology*, ed. Janet M. Box-Steffensmeier, Henry E. Brady and David Collier. New York: Oxford University Press.

- Buşoniu, Lucian, Alessandro Lazaric, Mohammad Ghavamzadeh, Rémi Munos, Robert Babuška and Bart De Schutter. 2012. Least-Squares Methods for Policy Iteration. In *Reinforcement Learning: State-of-the-Art*, ed. Marco Wiering and Martijn van Otterlo. New York: Springer.
- Buşoniu, Lucian, Robert Babuška, Bart De Schutter and Damien Ernst. 2010. *Reinforcement Learning and Dynamic Programming Using Function Approximators*. Boca Raton, FL: CRC Press.
- Bullock, John G., Donald P. Green and Shang E. Ha. 2010. “Yes, But What’s the Mechanism? (Don’t Expect an Easy Answer).” *Journal of Personality and Social Psychology* 98(4):550–558.
- Cameron, A. Colin and Pravin K. Trivedi. 2009. *Microeconometrics with Stata*. College Station, TX: Stata Press.
- Cameron, A. Colin and Pravin K. Trivedi. 2013. *Regression Analysis of Count Data*. 2nd ed. New York: Cambridge University Press.
- Chakraborty, Bibhas and Erica E.M. Moodie. 2013. *Statistical Methods for Dynamic Treatment Regimes: Reinforcement Learning, Causal Inference, and Personalized Medicine*. New York: Springer.
- Chang, Eric C.C., Mark Andreas Kayser, Drew A. Linzer and Ronald Rogowski. 2010. *Electoral Systems and the Balance of Consumer-Producer Power*. New York: Cambridge University Press.

- Chapman, Audrey R. and Patrick Ball. 2001. "The Truth of Truth Commissions: Comparative Lessons from Haiti, South Africa, and Guatemala." *Human Rights Quarterly* 23:1–23.
- Chen, Ding-Geng (Din) and Karl E. Peace. 2011. *Clinical Trial Data Analysis Using R*. Boca Raton, FL: CRC Press.
- Chow, Shein-Chung and Mark Chang. 2006. *Adaptive Design Methods in Clinical Trials*. Boca Raton, FL: Chapman & Hall/CRC.
- Cingranelli, David L. and David L. Richards. 2010. "The Cingranelli and Richards (CIRI) Human Rights Data Project." *Human Rights Quarterly* 32(2):401–424.
- Cingrenelli, David L. and David L. Richards. 1999. "Measuring the Level, Pattern and Sequence of Government Respect for Physical Integrity Rights." *International Studies Quarterly* 43(2):407–18.
- Cole, Stephen R. and Miguel Hernán. 2008. "Constructing Inverse Probability Weights for Marginal Structural Models." *American Journal of Epidemiology* 168(6):656–664.
- Collier, David, Fernando Daniel Hidalgo and Andra Olivia Maciuceanu. 2006. "Essentially Contested Concepts: Debates and Applications." *Journal of Political Ideologies* 11(3):211–246.
- Connolly, William. 1993. *The Terms of Political Discourse*. 3rd ed. Oxford, UK: Blackwell.

- Cox, D.R. 1958. *Planning of Experiments*. New York: Wiley.
- Curtis, Lesley H., Bradley G. Hammill, Eric Eisenstein, Judith M. Kramer and Kevin J. Amstrom. 2007. “Using Inverse Probability-Weighted Estimator in Comparative Effectiveness Analyses With Observational Database.” *Medical Care* 45(10):S103–S107.
- de Greiff, Pablo, ed. 2006. *The Handbook of Reparations*. Oxford, UK: Oxford University Press.
- de Greiff, Pablo and Roger Duthie, eds. 2009. *Transitional Justice and Development*. New York: Social Science Research Council.
- Dockner, Engelbert J., Steffen Jorgensen, Ngo Van Long and Gerhard Sorger. 2001. *Differential Games in Economics and Management Science*. New York: Cambridge University Press.
- Edwards, David. 2000. *Introduction to Graphical Modeling*. 2nd ed. New York: Springer.
- Egger, Peter H. and Maximilian von Ehrlich. 2013. “Generalized Propensity Scores for Multiple Continuous Variables.” *Economic Letters* 119(1):32–34.
- Ehring, Douglas. 1997. *Causation and Persistence: A Theory of Causation*. New York: Oxford University Press.
- Elster, Jon. 1978. *Logic and Society: Contradictions and Possible Worlds*. Chichester, UK: John Wiley & Sons.

- Elster, Jon. 1998. "Coming to Terms with the Past: A Framework for the Study of Justice in the Transition to Democracy." *European Journal of Sociology* 39(1):7–48.
- Elster, Jon. 2004. *Closing the Books: Transitional Justice in Historical Perspective*. Cambridge: Cambridge University Press.
- Elster, Jon, ed. 1986. *Rational Choice*. Oxford: Basil Blackwell.
- Elster, Jon, ed. 2006. *Retribution and Reparation in the Transition to Democracy*. New York: Cambridge University Press.
- Ertefaie, Ashkan and David A. Stephens. 2010. "Comparing Approaches to Causal Inference for Longitudinal Data: Inverse Probability Weighting versus Propensity Scores." *The International Journal of Biostatistics* 6(2):1–24.
- Fearon, James D. 1991. "Counterfactuals and Hypothesis Testing in Political Science." *World Politics* 43(2):169–95.
- Finnemore, Martha and Kathryn Sikkink. 1998. "International Norm Dynamics and Political Change." *International Organization* 52(4):887–917.
- Freeman, Mark. 2006. *Truth Commissions and Procedural Fairness*. New York: Cambridge University Press.
- Gallie, W. B. 1956. "Essentially Contested Concepts." *Proceedings of the Aristotelian Society* 56.
- Gelman, Andrew and Jennifer Hill. 2006. *Data Analysis Using Regression and Multilevel/Hierarchical Models*. New York: Cambridge University Press.



- Gibson, James L. 2009. *Overcoming Historical Injustices: Land Reconciliation in South Africa*. New York: Cambridge University Press.
- Glynn, Adam and Kevin M. Quinn. 2010. "An Introduction to the Augmented Inverse Propensity Weighted Estimator." *Political Analysis* 18(1):36–56.
- Glynn, Adam N. and Kevin M. Quinn. 2011. "Why Process Matters for Causal Inference." *Political Analysis* 19(3):273–286.
- Grimmer, Justin. 2013. "Text as Data." *Political Analysis* 21(3):267–97.
- Guo, Shenyang and Mark W. Fraser. 2010. *Propensity Score Analysis: Statistical Methods and Applications*. Thousand Oaks, CA: Sage.
- Haunschmied, Josef, Vladimir M. Veliov and Stafan Wrzaczek, eds. 2014. *Dynamic Game in Economics*. New York: Springer.
- Haurie, Alain and Georges Zaccour, eds. 2005. *Dynamic Games: Theory and Applications*. New York: Springer.
- Hayner, Priscilla. 2001. *Unspeakable Truths: Confronting State Terror and Atrocity*. New York: Routledge.
- Hayner, Priscilla. 2010. *Unspeakable Truths: Transitional Justice and the Challenge of Truth Commissions*. 2nd ed. New York: Routledge.
- Hazan, Pierre. 2006. "Measuring the Impact of Punishment and Forgiveness: A Framework for Evaluating Transitional Justice." *International Review of the Red Cross* 88(861):19–47.

- Heckman, James J. 2005. "The Scientific Model of Causality." *Sociological Methodology* 35:1–97.
- Hernán, Miguel A., Emilie Lanoy, Dominique Costagliola and James M. Robins. 2006. "Comparison of Dynamic Treatment Regimes via Inverse Probability Weighting." *Basic & Clinical Pharmacology & Toxicology* 98(3):237–242.
- Hirano, Keisuke and Guido W. Imbens. 2004. The Propensity Score with Continuous Treatments. In *Applied Bayesian Modeling and Causal Inference from Incomplete-Data perspectives*, ed. Andrew Gelman and Xiao-Li Meng. West Sussex, England: John Wiley & Sons.
- Ho, Daniel, Kosuke Imai, Gary King and Elizabeth Stuart. 2007. "Matching as Non-parametric Preprocessing for Reducing Model Dependence in Parametric Causal Inference." *Political Analysis* 15(3):199–236.
- Holland, Paul W. 1986. "Statistics and Causal Inference." *Journal of the American Statistical Association* 81(396):945–960.
- Holmes, William M. 2014. *Using Propensity Scores in Quasi-Experimental Designs*. Thousand Oaks, CA: Sage.
- Hong, Guanglei and Stephen W. Raudenbush. 2013. Heterogenous Agents, Social Interactions, and Causal Inference. In *Handbook of Causal Analysis for Social Research*, ed. Stephen L. Morgan. New York: Springer.
- Hoover, Kevin D. 2001. *Causality in Macroeconomics*. New York: Cambridge University Press.

- Hopkins, Daniel and Gary King. 2010. "A Method of Automated Nonparametric Content Analysis for Social Science." *American Journal of Political Science* 54(1):229–247.
- Horvitz, D.G. and D.J. Thompson. 1952. "A Generalization of Sampling Without Replacement From a Finite Universe." *Journal of the American Statistical Association* 47(260):663–685.
- Huntington, Samuel P. 1991. *The Third Wave: Democratization in the Late Twentieth Century*. Norman, OK: The University of Oklahoma Press.
- Imai, Kosuke and David A. van Dyk. 2004. "Causal Inference With General Treatment Regimes: Generalizing the Propensity Score." *Journal of the American Statistical Association* 99(467):854–866.
- Imbens, Guido W. 2000. "The Role of the Propensity Score in Estimating Dose-Response Functions." *Biometrika* 87(3):706–710.
- Inglehart, Ronald and Christian Welzel. 2005. *Modernization, Cultural Change, and Democracy: The Human Development Sequence*. New York: Cambridge University Press.
- IRIN. 2006. "Justice for a Lawless World? Rights and Reconciliation in a New Era of International Law." <http://www.irinnews.org/InDepthMain.aspx?InDepthId=7&ReportId=59468>. International Regional Information Networks.

- Isaacs, Rufus. 1965. *Differential Games: A Mathematical Theory with Applications to Warfare and Pursuit, Control and Optimization*. New York: John Wiley & Sons.
- Jackson, John E. and Ken Kollman. 2010. "A Formulation of Path Dependence with an Empirical Example." *Quarterly Journal of Political Science* 5(3):257–289.
- Joffe, Marshall M. and Paul R. Rosenbaum. 1999. "Invited Comment: Propensity Scores." *American Journal of Epidemiology* 159(4):327–339.
- Keele, Luke and Nathan J. Kelly. 2006. "Dynamic Models for Dynamic Theories: The Ins and Outs of Lagged Dependent Variables." *Political Analysis* 14(2):186–205.
- Kim, Hunjoon and Kathryn Sikkink. 2010. "Explaining the Deterrence Effect of Human Rights Prosecutions for Transitional Countries." *International Studies Quarterly* 54(4):939–963.
- Koller, Daphne and Nir Friedman. 2009. *Probabilistic Graphical Models: Principles and Techniques*. Cambridge, MA: The MIT Press.
- Kritz, Neil. 2004. Dealing with the Legacy of Past Abuses. In *Dealing with the Past: Critical Issues, Lessons Learned, and Challenges for Future Swiss Policy*, ed. M<sup>o</sup> Bleeker and Jonathan Sisson. Bern, Switzerland: Swisspeace.
- Kritz, Neil J. 1996. "Coming to Terms with Atrocities: A Review of Accountability Mechanisms for Mass Violation of Human Rights." *Law and Contemporary Problems* 59(4):127–152.

- Kritz, Neil J., ed. 1995a. *Transitional Justice: How Emerging Democracies Reckon with Former Regimes, Volume I: General Considerations*. Washington, D.C.: United States Institute of Peace Press.
- Kritz, Neil J., ed. 1995b. *Transitional Justice: How Emerging Democracies Reckon with Former Regimes, Volume II: Country Studies*. Washington, D.C.: United States Institute of Peace Press.
- Kritz, Neil J., ed. 1995c. *Transitional Justice: How Emerging Democracies Reckon with Former Regimes, Volume III: Laws, Rulings and Reports*. Washington, D.C.: United States Institute of Peace Press.
- Landman, Todd. 2006. *Studying Human Rights*. New York: Routledge.
- Lauritzen, Steffan L. 1996. *Graphical Models*. Oxford, UK: Clarendon Press.
- Lechner, Michael. 2004. Sequential Matching Estimation. Discussion Paper, 1042, Forschungsinstitut zur Zukunft der Arbeit.
- Lechner, Michael. 2006. Matching Estimating of Dynamic Treatment Models: Some Practical Issues. Discussion Paper, 2006-03, Universität St.Gallen.
- Lewis, David. 1973. *Counterfactuals*. Cambridge, MA: Harvard University Press.
- Lindberg, Steffan L., ed. 2009. *Democratization by Elections*. Baltimore, MD: The Johns Hopkins University Press.

- Lipset, Seymour Martin. 1959. "Some Social Requisites of Democracy: Economic Development and Political Legitimacy." *American Political Science Review* 53(1):69–105.
- Lohr, Sharon L. 2009. *Sampling: Design and Analysis*. 2nd ed. Boston, MA: Cengage Learning.
- Long, Nogo Van. 2010. *A Survey of Dynamic Games in Economics*. Singapore: World Scientific.
- Lu, Bo, Elaine Zanutto, Robert C. Hornik and Paul R. Rosenbaum. 2001. "Matching with Doses in an Observational Study of a Media Campaign Against Drug Abuse." *Journal of the American Statistical Association* 96(456):1245–1253.
- Luce, R. Duncan and Howard Raiffa. 1957. *Games and Decisions*. New York: John Wiley & Sons.
- Lunceford, Jared K. and Marie Davidian. 2004. "Stratification and Weighting Via the Propensity Score in Estimation of Causal Treatment Effects: A Comparative Study." *Statistics in Medicine* 23(19):2037–2060.
- Malamud-Goti, Jaime. 1996. *Game without End: State Terror and Politics of Justice*. Norman, OK: University of Oklahoma Press.
- Manski, Charles F. 2007. *Identification for Prediction and Decision*. Cambridge, MA: Harvard University Press.

- McFaul, Michael. 2002. "The Fourth Wave of Democracy and Dictatorship: Noncooperative Transitions in the Postcommunist World." *World Politics* 54(2):212–244.
- Mendeloff, David. 2004. "Truth-Seeking, Truth-Telling, and Postconflict Peacebuilding: Curb the Enthusiasm?" *International Studies Review* 6(3):355–80.
- Minow, Martha. 1998. *Between Vengeance and Forgiveness: Facing History after Genocide and Mass Violence*. Boston, MA: Beacon Press.
- Moodie, Erica E.M., Thomas S. Richardson and David A. Stephens. 2007. "Demystifying Optimal Dynamic Treatment Regimes." *Biometrics* 63(2):447–455.
- Moore, J.M. 1975. *Aristotle and Xenophon on Democracy and Oligarchy*. Berkeley, CA: The University of California Press.
- Morgan, Stephen L. and Christopher Winship. 2007. *Counterfactuals and Causal Inference: Methods and Principles for Social Research*. New York: Cambridge University Press.
- Morton, Rebecca B. and Kenneth C. Williams. 2010. *Experimental Political Science and the Study of Causality: From Nature to the Lab*. New York: Cambridge University Press.
- Mosley, Layna and Saika Uno. 2007. "Reeling to the Bottom or Climbing to the Top? Economic Globalization and Collective Labor Rights." *Comparative Political Studies* 40(8):923–48.

- Munck, Gerardo L. and Carol Skalnik Leff. 1997. "Modes of Transition and Democratization: South America and Eastern Europe in Comparative Perspective." *Comparative Politics* 29(3):343–362.
- Murphy, Susan. 2003. "Optimal Dynamic Treatment Regimes." *Journal of Royal Statistical Society: Series B* 65(2):331–366.
- Nahum-Shani, Inbal, Min Quian, Daniel Amirall, William E. Pelham, Beth Gnagy, Greg Fabino, Jim Waxmonkey, Jihnhee Yu and Susan Murphy. 2012. "Q-learning: A Data Analysis Method for Constructing Adaptive Interventions." *Psychological Methods* 17(4):478–494.
- Nickell, Stephen. 1981. "Biases in Dynamic Models with Fixed Effects." *Econometrica* 49(6):1417–1426.
- Nino, Carlos. 1996. *Radical Evils on Trial*. New Haven, CT: Yale University Press.
- North, Douglass C. 1990. *Institutions, Institutional Change and Economic Performance*. New York: Cambridge University Press.
- O'Donnell, Guillermo and Philippe Schmitter. 1986. *Transitions from Authoritarian Rule: Tentative Conclusions about Uncertain Democracies*. Baltimore, MD: The Johns Hopkins University Press.
- Olsen, Tricia D., Leigh A. Payne and Andrew G. Reiter. 2010. *Transitional Justice in Balance: Comparing Processes, Weighing Efficacy*. Washington, DC: United States Institute of Peace Press.



- Page, Scott E. 2006. "Path Dependence." *Quarterly Journal of Political Science* 1(1):87–115.
- Parmigiani, Giovanni and Lurdes Inoue. 2009. *Decision Theory: Principles and Approaches*. Chichester, UK: John Wiley & Sons.
- Pearl, Judea. 2000. *Causality: Models, Reasoning, and Inference*. New York: Cambridge University Press.
- Pearl, Judea. 2009. *Causality: Models, Reasoning, and Inference*. 2nd ed. New York: Cambridge University Press.
- Piantadosi, Steven. 2005. *Clinical Trials: A Methodological Perspective*. 2nd ed. Hoboken, NJ: Wiley.
- Pierson, Paul. 2004. *Politics in Time: History, Institutions, and Social Analysis*. Princeton, NJ: Princeton University Press.
- Pion-Berlin, Douglas. 1994. "To Prosecute or to Pardon? Human Rights Decision in the Latin American Southern Cone." *Human Rights Quarterly* 16(1):105–130.
- Poe, Steven C. and C. Neal Tate. 1994. "Repression of Human Rights to Personal Integrity in the 1980s: A Global Analysis." *American Political Science Review* 88(4):85–872.
- Poe, Steven C., Neal Tate and Linda Camp Keith. 1999. "Repression of Human Right to Personal Integrity Revisited: A Global Cross-National Study Covering the Years 1976-1993." *International Studies Quarterly* 43(2):291–313.

- Posner, Eric A. and Adrian Vermeule. 2004. “Transitional Justice as Ordinary Justice.” *Harvard Law Review* 117(3):761–825.
- Powell, Warren B. 2007. *Approximate Dynamic Programming: Solving the Curses of Dimensionality*. Hoboken, NJ: Wiley-Interscience.
- Przeworski, Adam, Jose Antonio Cheibub, Michael Alvarez and Fernando Limongi. 2000. *Democracy and Development: Political Institutions and Material Well-being in the World, 1950-1990*. New York: Cambridge University Press.
- Quinn, Kevin M., Burt L. Monroe, Michael Colaresi, Michael H. Crespin and Dragomir R. Radev. 2010. “How to Analyze Political Attention with Minimal Assumptions and Costs.” *American Journal of Political Science* 54(1):209–22.
- Robins, James. 2004. Optimal Structural Nested Models for Optimal Sequential Decisions. In *Proceedings of the Second Seattle Symposium on Biostatistics*, ed. Danyu Lin and Patrik J. Heagerty. New York: Springer.
- Robins, James, Donald Blevins, Grant Ritter and Michael Wulfsohn. 1992. “G-Estimation of the Effect of Prophylaxis Therapy for *Pneumocystis Carinii* Pneumonia on the Survival of AIDS Patients.” *Epidemiology* 3(4):319–336.
- Robins, James, Liliana Orellana and Andrea Rotnitzky. 2008. “Estimation and Extrapolation of Optimal Treatment and Testing Strategies.” *Statistics in Medicine* 27:4678–4721.

- Robins, James M. 1986. "A New Approach to Causal Inference in Mortality Studies with a Sustained Exposure Period - Application to Control of the Healthy Worker Survivor Effect." *Mathematical Modeling* 7(9-12):1393–1512.
- Robins, James M., Andrea Rotnitzky and Lue Ping Zhao. 1994. "Estimation of Regression Coefficients When Some Regressors Are Not Always Observed." *Journal of the American Statistical Association* 89(427):846–866.
- Rogowski, Ronald and Mark Andreas Kayser. 2002. "Majoritarian Electoral Systems and Consumer Power: Price-Level Evidence from the OECD Countries." *American Journal of Political Science* 46(3):526–539.
- Roht-Arriaza, Naomi and Javier Mariezcurrena, eds. 2006. *Transitional Justice in the Twenty-First Century: Beyond Truth Versus Justice*. New York: Cambridge University Press.
- Rosenbaum, Paul R. 2002. *Observational Studies*. 2nd ed. New York: Springer.
- Rosenbaum, Paul R. 2010. *Design of Observational Studies*. New York: Springer.
- Rosenbaum, Paul R. and Donald Rubin. 1983. "The Central Role of the Propensity Score in Observational Studies for Causal Effects." *Biometrika* 70(1):41–55.
- Rubin, Donald B. 1980. "Comment on "Randomization Analysis of Experimental Data in the Fisher Randomization Test" by Basu." *Journal of the American Statistical Association* 75(371):591–593.

- Rubin, Donald B. 2006. *Matched Sampling for Causal Effects*. New York: Cambridge University Press.
- Schotter, Andrew. 1981. *The Economic Theory of Social Institutions*. Cambridge, UK: Cambridge University Press.
- Schwartz, Howard M. 2014. *Multi-Agent Machine Learning: A Reinforcement Approach*. Hoboken, NJ: Wiley.
- Sekhon, Jasjeet S. 2004. “Quality Meets Quantity: Case Studies, Conditional Probability, and Counterfactuals.” *Perspectives on Politics* 2(2):281–293.
- Sekhon, Jasjeet S. 2008. The Neyman-Rubin Model of Causal Inference and Estimation via Matching Methods. In *The Oxford Handbook of Political Methodology*, ed. Janet M. Box-Steffensmeier, Henry E. Brady and David Collier. Oxford University Press.
- Shaw, Rosalind and Lars Waldorf. 2010. Introduction: Localizing Transitional Justice. In *Localizing Transitional Justice: Interventions and Violence after Mass Violence*, ed. Rosalind Shaw, Lars Waldorf and Pierre Hazan. Stanford, CA: Stanford University Press.
- Shaw, Rosalind, Lars Waldorf and Pierre Hazan, eds. 2010. *Localizing Transitional justice: Interventions and Violence after Mass Violence*. Stanford, CA: Stanford University Press.
- Sikkink, Kathryn and Carrie Booth Walling. 2007. “The Impact of Human Rights Trials in Latin America.” *Journal of Peace Research* 44(4):427–445.

- Snyder, Jack and Leslie Vinjamura. 2003/4. "Trials and Errors: Principle and Pragmatism in Strategies of International Justice." *International Security* 28(3):5–44.
- Spirtes, Peter, Clark Glymour and Richard Scheines. 2000. *Causation, Prediction, and Search*. 2nd ed. Cambridge, MA: The MIT Press.
- Stan, Lavinia and Nadya Nedelsky, eds. 2013a. *Encyclopedia of Transitional Justice, vol. 1*. New York: Cambridge University Press.
- Stan, Lavinia and Nadya Nedelsky, eds. 2013b. *Encyclopedia of Transitional Justice, Vol. 2*. New York: Cambridge University Press.
- Stan, Lavinia and Nadya Nedelsky, eds. 2013c. *Encyclopedia of Transitional Justice, vol. 3*. New York: Cambridge University Press.
- Sutton, Richard S. and Andrew G. Barto. 1998. *Reinforcement Learning: An Introduction*. Cambridge, MA: The MIT Press.
- Szepesvari, Csaba. 2010. *Algorithms for Reinforcement Learning*. San Rafael, CA: Morgan and Claypool.
- Teitel, Ruti G. 2000. *Transitional Justice*. Oxford: Oxford University Press.
- Teitel, Ruti G. 2003. "Transitional Justice Genealogy." *Harvard Human Rights Journal* 16:69–94.
- Tepperman, Jonathan D. 2002. "Truth and Consequences." *Foreign Affairs* March/April(81):128–145.

- Tetlock, Philip E. and Aaron Belkin. 1996. Counterfactual Thought Experiments in World Politics: Logical, Methodological, and Psychological Perspectives. In *Counterfactual Thought Experiments in World Politics: Logical, Methodological, and Psychological Perspectives*, ed. Philip E. Tetlock and Aaron Belkin. Princeton: Princeton University Press.
- Theissen, Gunnar. 2008. Object of Trust and Hatred: Public Attitudes Toward the TRC. In *Truth and Reconciliation in South Africa; Did the TRC Deliver?*, ed. Audrey R. Chapman and Hugo van der Merwe. Philadelphia, PA: University of Pennsylvania Press.
- Thoms, Oskar N. T., James Ron and Roland Paris. 2010. "State-Level Effects of Transitional Justice: What Do We Know?" *The International Journal of Transitional Justice* 4(3):1–29.
- Thoms, Oskar N.T., James Ron and Roland Paris. 2008. "The Effects of Transitional Justice Mechanisms." Center for International Policy Studies, University of Ottawa.
- Villani, Cédric. 2009. *Optimal Transport: Old and New*. Berlin, Germany: Springer.
- Villani, Cédric, ed. 2003. *Topics in Optimal Transportation*. Providence, RI: American Mathematical Society.
- von Neumann, John and Oskar Morgenstern. 1980. *Theory of Games and Economic Behavior*. 3rd ed. Princeton, NJ: Princeton University Press.
- Waldron, Jeremy. 1992. "Superseding Historical Injustice." *Ethics* 103(1):4–28.

- Watkins, C.J.C.H. 1989. Learning from Delayed Rewards. PhD thesis , Cambridge University.
- Weber, Max. 1949. *The Methodology of the Social Sciences*. New York: Free Press.
- Wiebelhaus-Brahm, Eric. 2010. *Truth Commissions and Transitional Societies: The Impact on Human Rights and Democracy*. London: Routledge.
- Wiering, Marco and Martijn van Otterlo, eds. 2012. *Reinforcement Learning: State-of-the-Art*. New York: Springer.
- Winship, Christopher and Stephen Morgan. 1999. “The Estimation of Causal Effects From Observational Data.” *Annual Review of Sociology* 25:659–706.
- Wood, Reed M. and Mark Gibney. 2010. “The Political Terror Scale (PTS): A Re-introduction and a Comparison to CIRI.” *Human Rights Quarterly* 32(2):367–400.
- Woodward, James. 2003. *Making Things Happen: A Theory of Causal Explanation*. Oxford: Oxford University Press.
- Woodward, James F. 2009. Agency and Interventionist Theories. In *The Oxford Handbook of Causation*, ed. Helen Beebe, Christopher Hitchcock and Peter Menzies. New York: Oxford University Press.
- Zalaquett, José. 1992. “Balancing Ethical Imperatives and Political Constraints: The Dilemmas of New Democracies Confronting Past Human Rights Violations.” *Hastings Law Journal* 43:1425–38.