PROJECT MUSE®

Handling Limited Overlap in Observational Studies with Cardinality Matching

Giancarlo Visconti, José R. Zubizarreta

Observational Studies, Volume 4, Issue 1, 2018, pp. 217-249 (Article)

Published by University of Pennsylvania Press

DOI: <https://doi.org/10.1353/obs.2018.0012>

➔ *For additional information about this article*

<https://muse.jhu.edu/article/793377/summary>

Handling Limited Overlap in Observational Studies with Cardinality Matching

Giancarlo Visconti

*Department of Political Science
Purdue University
West Lafayette, IN 47907*

gviscont@purdue.edu

José R. Zubizarreta

*Department of Health Care Policy and Department of Statistics
Harvard University
Boston, MA 02115*

zubizarreta@hcp.med.harvard.edu

Abstract

A common problem encountered in observational studies is limited overlap in covariate distributions across treatment groups. To address this problem, and avoid strong modeling assumptions, it has become common practice to restrict analyses to the portions of the treatment groups that overlap or, ultimately, are balanced in their covariate distributions. Often, this is done by matching on the estimated propensity score or coarsened versions of the observed covariates. A recent alternative methodology that, in a sense, encompasses these two approaches is cardinality matching. Cardinality matching is a flexible matching method that uses integer programming to find the largest matched sample that is balanced according to criteria specified before matching by the investigator. In this paper, we apply and illustrate the method of cardinality matching and show how to use it to directly balance several features of the covariates, including their trajectories in time and their distributions, without requiring exact matching. We demonstrate how cardinality matching addresses the problem of limited overlap using the original covariates, as opposed to a summarized or coarsened version of them. We discuss how this method can be extended to build matched samples that are not only balanced but also representative of a target population by design. We also show how this method enhances sensitivity analyses for hidden biases. We explain these advancements through an observational study of the electoral impact of the 2010 earthquake in Chile.

Keywords: Causal inference, Matched sampling, Natural disasters, Natural experiments, Observational studies

1. Introduction

A guiding principle in the design of observational studies is to approximate the ideal study that would be conducted if it were possible to do it by controlled experimentation (Dorn 1953; Cochran and Rubin 1973; Rosenbaum 2010). In an experiment, randomization tends to produce treatments groups that are comparable in terms of both observed and unobserved covariates and provides a basis for valid statistical inference that does not rely on

distributional assumptions (see, e.g., Armitage 1982). In contrast, in observational studies randomization is absent, the treatment assignment mechanism is unknown, and there may be systematic covariate differences between the treatment groups that can bias effect estimates. In observational studies, matching methods are often used to adjust for differences in observed covariates, thus emulating the structure of the ideal randomized study and facilitating sensitivity analyses of the outcome results to biases due to unobserved covariates (see, e.g., Chapter 4 of Rosenbaum 2002).

In practice, a common problem encountered in observational studies is the lack of common support or the limited overlap of the covariate distributions across treatment groups. As discussed by Ho et al. (2007) and Crump et al. (2009), limited overlap can lead to covariate imbalances and produce estimates that are overly sensitive to model misspecification. To address this problem, it is common practice with matching methods to “trim” the treatment and control samples, and restrict the analyses to the portions of the samples that overlap or, ultimately, that are balanced. Often, this is done using a summary measure of the covariates such as the propensity score or the Mahalanobis distance (e.g., Austin 2011), although alternative matching methods that select a subset of the observations using the original observed covariates instead of a summary of them have also been proposed (see Fogarty et al. 2016). Perhaps the most popular of these methods is Coarsened Exact Matching, or CEM (Iacus et al. 2012). In short, this method first coarsens or discretizes the observed covariates, then forms multidimensional strata using the coarsened covariate values, and finally matches the treated and control units in the same stratum, discarding those units with no treated or control counterpart for each stratum.

While matching on the estimated propensity score or coarsened versions of the observed covariates are useful and popular methods, they can fail to balance the original covariates (Diamond and Sekhon 2013) or use very few observations (King et al. 2016). A recent alternative method that overcomes these limitations is cardinality matching (Zubizarreta et al. 2014).

Taking advantage of recent advancements in mathematical programming (particularly, in integer programming), cardinality matching maximizes the cardinality or size of the matched sample that satisfies the investigators’ requirements for covariate balance. As discussed in Zubizarreta (2012), these requirements are flexible and can enforce balance for different features of the empirical distributions of the observed covariates, such as moments, marginals and, in principle, their entire joints. These requirements can also enforce exact matching for coarsened covariates, as in CEM. As we argue in this paper, this flexibility in balancing covariates can be very important for avoiding the curse of dimensionality in exact matching with several covariates and facilitate the use of more observations in the analysis. In this paper, we describe and apply the method of cardinality matching in a study about natural disasters and electoral outcomes.

What are the benefits of cardinality matching? First, it handles limited overlap in covariate distributions by using the original covariates as opposed to summaries or coarsened versions of them, directly balancing the covariates, and finding the maximum number of observations for any given covariate balance criteria. Importantly, with cardinality matching these criteria are flexible: one can require different forms of balance covariates from mean balance

to exact matching. Thus, in a sense, subject to a matching structure (e.g., a one-to-one matching structure) CEM is a particular case of cardinality matching, where all the balance constraints enforce exact matching for coarsened covariates. In a study like ours, which examines the impact of an earthquake on electoral outcomes at the county level, where there are not many county-level observations to begin with, the flexibility of cardinality matching allows us to use more observations from the data than other matching methods for any given covariate balance criteria.

The rest of this paper is organized as follows. In Section 2, we review cardinality matching, explain its mathematical programming formulation, and discuss different forms of covariate balance that can be enforced with it. In Section 3, we describe elections in Chile, the 2010 Chilean earthquake, our longitudinal county data, and the study design. In Section 4, we evaluate the matched sample and report effect estimates. In Section 5, we assess the sensitivity and generalizability of our findings. In Section 6, we compare cardinality matching to other matching techniques. Finally, in Section 7, we close with a summary of the study.

2. Maximizing the cardinality of a balanced matched sample

2.1 Review of cardinality matching

The goal of matching in observational studies is to remove the bias in effect estimates due to imbalances in the observed covariates (Rosenbaum 2015). Though this goal would ideally be accomplished using all of the available treated and control observations, this is rarely possible in practice because of limited overlap in their covariate distributions. In view of this limitation, investigators often subset or trim the treated and control samples and confine their analyses to subsamples for which covariates are balanced.

Perhaps the most common approach to this problem relies on the propensity score, whereby treated (control) units with an estimated propensity score outside the range of the control (treated) units are discarded (Stuart 2010). Afterwards, treated and control units are matched (for example, using nearest neighbor matching (Rubin 1973; Abadie et al. 2004) or optimal matching (Rosenbaum 1989; Hansen 2007)) and then covariate balance is checked. These steps are summarized in Algorithm 1 below. As discussed by Hainmueller (2012) and Diamond and Sekhon (2013), this procedure typically requires many iterations in order to achieve covariate balance.

This procedure can be improved by using the covariate balancing propensity score (Imai and Ratkovic 2014), a propensity score model that penalizes fits for which covariates are imbalanced. Other methods for selecting subsamples of treated and control observations are proposed by Crump et al. (2009) and Rosenbaum (2012). These approaches, however, address the problem of limited overlap using a one-dimensional summary of the covariates rather than the original covariates themselves.

An alternative matching method, coarsened exact matching (CEM; Iacus et al. 2012), selects a subset of the observations but using the original covariates. This method first coarsens

the individual covariates (for example, by transforming its values into quintile or decile indicators), then exactly matches the treated and control units according to the coarsened covariates, and discards units that cannot be exactly matched. While exact matching is the ideal form of covariate balance, it suffers from the curse of dimensionality, as in actual practice the number of exact matching categories tends to explode combinatorially. In consequence, exact matching in practice tends to be relegated to a few covariates of overriding importance (see, for example, Chapter 9 of Rosenbaum 2010).

Another recent matching method that selects a subset of the treated and control units using the original covariates is cardinality matching (Zubizarreta et al. 2014). This method solves a linear integer programming problem to maximize the cardinality or size of a matched sample subject to flexible constraints on covariate balance. In their most stringent form, these constraints can require exact matching on the observed covariates, but they can also require other weaker forms of covariate balance: aggregate balance of low-dimensional joint distributions via strength- k balancing (Hsu et al. 2015), balance of marginal distributions by balancing a coarsened version of the Kolmogorov-Smirnov statistic (Zubizarreta 2012) or by fine balancing (Rosenbaum et al. 2007), and balance of moments by mean balancing suitable transformations of the covariates (see Zubizarreta 2012). In this manner, cardinality matching subsets or trims the treatment and control samples and balances multiple covariates in one step, and thus finds the largest matched sample that is balanced.

As described in Section 2.2, cardinality matching finds the largest matched sample that is balanced by solving a constrained linear integer programming problem where the constraints enforce a given matching structure (in our case study, one-to-one or pair matching) and pre-specified forms of covariate balance (ranging from mean balance to exact matching and including different forms of aggregate distributional balance) and the objective function maximizes the total number of matched pairs. Cardinality matching guarantees the optimal solution (or, in some cases, a near-optimal solution within a provable optimality gap) by exploiting recent advancements in algorithms and computation, including advanced heuristics, presolve methods, linear programming, and faster compilers (see Gurobi 2017 for details). In short, cardinality matching directly targets bias removal via flexible covariate balance constraints and variance reduction via an objective function on sample size by solving an optimization problem as opposed to greedy or heuristic methods.

Algorithms 1 and 2 show the basic steps involved in (i) standard matching methods based on the estimated propensity score and (ii) cardinality matching. With standard matching methods, steps 2-5 typically require many repetitions in order to satisfy the covariate balance requirements, whereas with cardinality matching the covariate balance requirements are satisfied in one step without needing to estimate the propensity score.

In cardinality matching, finding the largest matched sample that satisfies the covariate balance requirements is often followed by rematching the balanced matched sample in order to minimize the covariate distances between matched units. Since the rematched sample is composed of the same treated and control units as the balanced matched sample, the aggregate balance measures are preserved after rematching. However, if the covariates used in calculating the distances are strong predictors of the outcome, then this rematching will reduce the heterogeneity in the matched pair differences in the outcomes. As discussed in

Rosenbaum (2005) and Zubizarreta et al. (2014), this translates into greater efficiency and lower sensitivity to unobserved covariates. In Section 5, we illustrate these gains in the context of our case study of the effect of the Chilean earthquake and subsequent reconstruction process on electoral outcomes.

Algorithm 1 Matching with standard matching methods.

0. Specify the covariate balance requirements (e.g., mean balance).

Repeat:

1. Estimate the propensity score or another summary of the covariates.
2. Trim extreme observations according to the summary measure.
3. Match on the summary measure (e.g., using nearest neighbor matching).
4. Assess covariate balance.

Until:

The matched sample satisfies the covariate balance requirements.

Algorithm 2 Matching with cardinality matching.

0. Specify the covariate balance requirements (e.g., mean balance).
 1. Find the largest matched sample that satisfies the covariate balance requirements.
 2. Rematch the matched sample to minimize covariate distances between matched units.
-

2.2 Mathematical programming formulation

Let $\mathcal{T} = \{1, \dots, T\}$ and $\mathcal{C} = \{1, \dots, C\}$ be the sets of treated and control units available before matching, with $T \leq C$. Define $t \in \mathcal{T}$ and $c \in \mathcal{C}$ as the corresponding indexes of these two sets. Write p to index the observed covariates in $\mathcal{P} = \{1, \dots, P\}$, so that $\mathbf{x}_t = \{x_{t1}, \dots, x_{tp}\}$ and $\mathbf{x}_c = \{x_{c1}, \dots, x_{cp}\}$ are the vectors of observed covariates for treated unit t and control unit c , respectively.

In its simplest form, the goal of cardinality matching is to find the largest pair-matched sample that is balanced. Using the binary decision variable $m_{tc} = 1$ if treated unit t is matched to control unit c , and $m_{tc} = 0$ otherwise, we can express this goal as the following objective function (1) subject to the following matching and balancing constraints (2) and (3), respectively. We wish to maximize the size of the matched sample,

$$\underset{\mathbf{m}}{\text{maximize}} \quad \sum_{t \in \mathcal{T}} \sum_{c \in \mathcal{C}} m_{tc}, \quad (1)$$

subject to matching constraints that enforce forming pairs,

$$\begin{aligned} \sum_{c \in \mathcal{C}} m_{tc} &\leq 1, \quad t \in \mathcal{T}, \\ \sum_{t \in \mathcal{T}} m_{tc} &\leq 1, \quad c \in \mathcal{C}, \end{aligned} \quad (2)$$

and balancing constraints that enforce the pair-matched samples to have similar covariate distributions

$$\left| \sum_{t \in \mathcal{T}} \sum_{c \in \mathcal{C}} m_{tc} (f_k(x_{tp}) - f_k(x_{cp})) \right| \leq \varepsilon_p \sum_{t \in \mathcal{T}} \sum_{c \in \mathcal{C}} m_{tc}. \quad (3)$$

In (3), f_k is a suitable function that transforms the observed covariates for each balance condition $k \in \mathcal{K}$, and $\varepsilon_p \geq 0$ is a tolerance, both of which are defined by the investigator. In Section 2.3 below, we discuss different forms of covariate balance, but before a few remarks are in order.

First, in cardinality matching, bias reduction takes precedence over variance reduction. (1)-(3) above maximize the size of the matched study, subject to removing biases due to imbalances in the observed covariates. This follows the recommendations of Cochran and Chambers (1965, Section 2.2) and Rosenbaum (2010, Section 8.7), who state that the main objective of matching in observational studies is to balance observed covariates, and that increasing efficiency is typically a second-order concern, as biases that do not decrease as the sample size increases tend to dominate the mean squared error in large samples (see also Haviland et al. 2007 and King et al. 2016).

Second, from a more formal standpoint, it is easy to prove that under a homoskedastic constant additive treatment effect outcome model, cardinality matching minimizes the variance of a difference-in-means effect estimator that is approximately unbiased according to the covariate balance constraints (see Kilcioglu et al. 2016).

Third, as discussed by Imbens (2015, page 382), when confronted with limited overlap in covariate distributions, “with multiple covariates, it is difficult to see what trimming would need to be done.” Often, the treated and control samples are trimmed using the propensity score, to then match the observations and check covariate balance on the original covariates (as in Algorithm 1). Cardinality matching essentially performs these three tasks in one step, “trimming” the treated and control samples using the original covariates.

Fourth, from a computational standpoint, cardinality matching solves an integer programming problem, so in the abstract cardinality matching is an NP-hard optimization problem. In practice, however, relatively large instances of this problem (with several hundreds and, in some instances, several thousands of observations) can be solved quickly by using modern solvers. Also, using approximation algorithms can further shorten the running time. For example, in the LaLonde (1986) data set (which has 185 and 15,992 treatment and control observations, and 8 covariates), using the function `cardmatch` in the `designmatch` package (Zubizarreta et al. 2017) for R with the option `gurobi` for the optimizer, one can find the largest pair-matched data set for which all the covariates are mean-balanced in less than one second. Of course, ultimately the performance of the method will depend on the data set at hand and the optimizer used.

Finally, in his influential article on observational studies, Cochran and Chambers (1965) give two basic pieces of advice for designing such studies (Rubin 2006, page 15): (i) “when selecting samples for study, make sure that they are large enough and have complete enough data to allow effects of practical importance to be estimated, and avoid treatment and control groups with large initial differences on confounding variables;” (ii) “use both the statistician

and the subject-matter expert in the planning stages.” Cardinality matching achieves the first goal because it finds the largest matched sample that is balanced, and facilitates the second goal because the knowledge of the subject-matter expert can be incorporated into the matching problem through the balance constraints (3). Again, these constraints are flexible and can require exactly matching all of the covariates or just balancing their means. We describe these constraints in what follows.

2.3 Forms of covariate balance

In our case study about the impact of the earthquake, we enforce different forms of covariate balance via the balance constraints (3). For some covariates we enforce mean balance. To do this, we let f be equal to the identity so that (3) bounds the differences in means between the matched treated and control groups

$$\left| \frac{\sum_{t \in \mathcal{T}} \sum_{c \in \mathcal{C}} m_{tc} x_{tp}}{\sum_{t \in \mathcal{T}} \sum_{c \in \mathcal{C}} m_{tc}} - \frac{\sum_{t \in \mathcal{T}} \sum_{c \in \mathcal{C}} m_{tc} x_{cp}}{\sum_{t \in \mathcal{T}} \sum_{c \in \mathcal{C}} m_{tc}} \right| \leq \varepsilon_p \quad (4)$$

and choose ε_p so that none of the standardized differences in means are greater than 0.1 (Rosenbaum and Rubin 1985).

For other covariates of greater prognostic importance, we also enforce distributional balance by modifying (3) in two ways. First, to approximately balance the marginal distribution of a continuous covariate, we define indicators for the quantiles of the covariate in the treatment group and then mean balance these indicators via (4). Effectively, this balances a coarsened version of the Kolmogorov-Smirnov statistic (Zubizarreta 2012). Second, to perfectly balance the marginal distribution of a nominal covariate — a form of balance called fine balance (Rosenbaum et al. 2007) — we let $\varepsilon_p = 0$ and $f(x_p) = \mathbb{1}_{\{x_p=b\}}$, where $b \in \mathcal{B}$ indexes the categories of the nominal covariate, so that (3) becomes

$$\sum_{t \in \mathcal{T}} \sum_{c \in \mathcal{C}} m_{tc} \mathbb{1}_{\{x_{tp}=b\}} = \sum_{t \in \mathcal{T}} \sum_{c \in \mathcal{C}} m_{tc} \mathbb{1}_{\{x_{cp}=b\}}, \quad b \in \mathcal{B}. \quad (5)$$

This constraint makes the counts of treated and control units in the matched sample be the same for each category of the nominal covariate, but without constraining units to be matched within the same category (as in exact matching).

In the context of (3), it is straightforward to enforce exact matching as follows

$$\sum_{t \in \mathcal{T}} \sum_{c \in \mathcal{C}} m_{tc} \mathbb{1}_{\{x_{tp}=b\} \cap \{x_{cp} \neq b\}} = 0, \quad b \in \mathcal{B}. \quad (6)$$

Of course, this is the ideal form of covariate balance, but also the most stringent one. See Zubizarreta (2012) and Zubizarreta et al. (2014) for a related discussion on forms of covariate balance.

2.4 Constructing a representative matched sample

The above matching framework can easily be extended to anchor the matched samples so that they are not only balanced, but also balanced around the distribution of a target population of policy interest (Bennett et al. 2018; Hirshberg and Zubizarreta 2018).

Let F be the joint cumulative distribution function of the P observed covariates of a target population of policy interest. Denote o_q^F as the q -th moment of F with $q \in \mathcal{Q} = \{1, \dots, Q\}$. For each moment $q \in \mathcal{Q}$, impose the following constraints to the matching problem

$$\left| \frac{\sum_{t \in \mathcal{T}} \sum_{c \in \mathcal{C}} m_{tc} \tilde{x}_{tq}}{\sum_{t \in \mathcal{T}} \sum_{c \in \mathcal{C}} m_{tc}} - o_q^F \right| \leq \epsilon_q \quad (7)$$

and

$$\left| \frac{\sum_{t \in \mathcal{T}} \sum_{c \in \mathcal{C}} m_{tc} \tilde{x}_{cq}}{\sum_{t \in \mathcal{T}} \sum_{c \in \mathcal{C}} m_{tc}} - o_q^F \right| \leq \epsilon_q \quad (8)$$

where $\tilde{x}_{\cdot q}$ is a suitable transformation of the observed covariates (depending on the moment of the target distribution to be balanced; Zubizarreta 2012) and ϵ_q is a scalar determined by the investigator. Under a MAR-type assumption (Little and Rubin 2014) — namely, that selection to the study sample is governed only by observed covariates, and that every unit in the target population has a positive probability of being selected into the study sample — plus strong ignorability (Rosenbaum and Rubin 1983), we can estimate conditional average treatment effects in the population. These ideas are implemented in the R package `designmatch` and constitute an alternative to regression estimation, whose estimates can have limited external validity (Aronow and Samii 2016). We illustrate how to use these ideas in Section 5.3.

3. Study of the impact of an earthquake on electoral outcomes

3.1 Earthquakes and Elections in Chile

After the end of the dictatorship in 1990, the opposition center-left coalition “Concertación” won four presidential elections in a row. In January 2010, Sebastián Piñera was elected president, representing the center-right coalition “Coalición por el Cambio.” In February 2010, a few days before Piñera’s inauguration, the country was shattered by an 8.8 magnitude earthquake. Piñera had to deal with the aftermath of the disaster for the duration of his administration, and the earthquake and its consequences were a central political issue for the next four years (Arana 2016).

The central-southern regions of Chile were severely affected by this natural disaster. The earthquake, and the tsunami that followed, killed over 500 people and injured about 12,000 more (Choi 2012). The disaster affected six of the fifteen regions of Chile.¹ The areas that experienced the highest-intensity shaking suffered significant damage to roads, ports,

¹ These are the regions of Valparaíso, O’Higgins, Maule, Biobío and Araucanía, and the Metropolitan Region.

bridges, airports, utilities, and communication networks (Hinrichs et al. 2011). About 220,000 homes, 3,700 schools, 17 hospitals, and 212 bridges were destroyed. The economic losses to the country were estimated at \$30 billion, or 18% of the country’s gross domestic product (McClean 2012).

People living in the most affected regions directly observed the reconstruction process, such as repairs to schools, hospitals, and roads, and the distribution of financial relief, for a long period of time. In November 2013, three years after the earthquake and mere weeks before the next presidential election, the government was still building new homes for residents in the affected areas,² and debate about the success of the reconstruction efforts were an important part of the campaign.³

Chile provides a unique opportunity to test the effects of earthquakes on electoral outcomes for three main reasons: (i) the fifteen regions of the country had suffered from earthquakes throughout the 20th century—events that cannot be anticipated; (ii) the country has a centralized political structure; and (iii) the country has low electoral volatility. The first and second points increase the comparability between counties affected and unaffected by the earthquake because all of them are vulnerable to earthquakes and have the same political configuration. Regarding the third point, the low electoral volatility in Chile means that previous electoral results (pretreatment measures of the outcomes) have strong predictive power over the subsequent electoral results (outcome of interest). The first post-earthquake presidential election provides the helpful conditions for understanding the effects of this particular natural disaster⁴ on voters’ electoral choices.⁵

3.2 Two competing arguments

What are the electoral consequences of natural disasters? The literature can be divided by two main arguments. The first posits that voters will always punish the government for negative shocks, including those whose origin is beyond the current administration’s responsibility and scope of action, such as droughts and floods (Achen and Bartels 2004, 2016). In particular, “incumbents will pay at the polls for bad times, even in situations where objective observers can find little rational basis to suppose that those incumbents have had any part in producing the voters’ pain” (Achen and Bartels 2004, page 7). Affected voters will simply blame the incumbent party “any time their own well-being falls below ‘normal’ levels, regardless of whether the incumbents have performed well or badly” (Achen and Bartels 2016, page 138). This explanation invokes citizens’ emotional distress and tendency to channel their frustration and misfortune into voting for the challenger. The second argument posits that voters may punish or reward the incumbent based on their response to the catastrophe (Healy and Malhotra 2010; Gasper and Reeves 2011). Specifically, “observing

² “Inauguran viviendas para damnificados del terremoto en Tomé.”, *La Nación*, November 13, 2013.

³ “Presidente defendió reconstrucción tras crítica de Bachelet.”, *La Nación*, August 30, 2013.

⁴ It is important to note that the treatment is not just the earthquake, but also the subsequent reconstruction process as well. The government’s provision of public and private goods and the evaluation of its performance are crucial variables for understanding the direction of the effects of the earthquake.

⁵ Carlin et al. (2014) study the effects of this earthquake on democratic legitimacy, but to our knowledge there have been no studies on its impact on electoral results.

that incumbents are adversely affected by natural disasters, does not necessarily mean that voters are irrational. Even though government cannot be blamed for the adverse natural events themselves, they can be held responsible for mitigation, response, and recovery” (Healy and Malhotra 2010, page 195). Therefore, this argument implies that voters have the capacity to judge authorities’ competence and that disaster victims do not irrationally punish incumbents after natural catastrophes.

In this paper, we focus on effect estimates and not on the causal mechanisms explaining the consequences of the disaster and the reconstruction process. The bulk of the literature has mainly studied the political consequences of floods (Bechtel and Hainmueller 2011; Visconti 2018), hurricanes (Sinclair et al. 2011; Chen 2013), tornadoes (Healy and Malhotra 2010), severe weather events (Gasper and Reeves 2011), forest fires (Lazarev et al. 2014), and volcanic eruptions (Maldonado et al. 2016). However, there is little evidence of the effects of earthquakes on electoral outcomes.

3.3 Longitudinal county data

We use data from all counties in Chile covering a time span of 10 years (1999-2009). This dataset provides county characteristics before the earthquake (i.e., the covariates), the intensity of this disaster (i.e., the treatment), and the electoral results after the earthquake (i.e., the outcome).

Because Chile has low levels of electoral volatility (Roberts 2013), the most important prognostic covariates are previous electoral results. We include the electoral results for the center-right, center-left, and alternative candidates in the last three presidential elections (1999, 2005, and 2009) using official data from the National Electoral Service. This allows us to have longitudinal county data of the pretreatment measures of the outcome variable.

Previous studies have shown that socioeconomic characteristics contribute to Chilean voters’ political preferences (Altman 2004; López 2004; Navia et al. 2008; Luna 2010; Calvo and Murillo 2012). We include the Human Development Index (HDI) calculated by the United Nations Development Programme at the county level in 2003⁶ and the poverty levels calculated the year before the earthquake by the Ministry of Social Development. Altman (2004) argues that the relationship between voting for the center-right coalition and the HDI follows a U-shape. This suggests we should balance the entire distribution of this covariate and not just its mean.

We also include a number of demographic characteristics at the county level from the 2002 national census that are important predictors of the vote.⁷ Urban voters, especially from larger cities, are more likely to vote for the center-left coalition, while voters from smaller cities or rural localities are more likely to vote for the center-right coalition (González 1999).

⁶ This index represents the average of three indicators: health, education, and income. The last time the UNDP computed these indices at the county level before the earthquake was in 2003.

⁷ This was the last census conducted before the earthquake.

For this reason, we also include the proportion of rural citizens and the total population at the county level.

3.4 Matched design

We use two criteria to define the exposed and control groups. The first criterion relies on the official reconstruction plan, in which the government identified six of the fifteen regions in the country as critically affected by the earthquake. Within these six affected regions, however, there was considerable variation in the intensity of the earthquake: the counties located farther from the epicenter or closer to the Andes were less affected. Furthermore, the official reconstruction plan shows that 87% of emergency houses were built in three of these six regions (Government of Chile 2010). Therefore, we need a second criterion to identify the most affected counties within the six regions.

To do this, we use the peak ground acceleration (PGA) as a measure of the intensity of the earthquake. Following Zubizarreta et al. (2013), in the exposed group we include the counties that had a PGA greater than 0.275 g. In the control group, we include the counties located in the nine unaffected regions, which were not part of the reconstruction plan. We also exclude the counties with a PGA lower than 0.275 g located in the affected regions, because extreme treatment conditions rather than marginal exposure tend to produce results less sensitive to hidden biases (Rosenbaum 2004). This use of extreme exposures is consistent with Zubizarreta et al. (2013), where they were used to study the effects of the 2010 earthquake on posttraumatic stress. As a way to check the suitability of these definitions, the counties placed in the control group had an average PGA of 0.06 g, whereas the counties placed in the exposed group had an average PGA of 0.29 g.

These criteria yield 77 counties in the exposed group and 95 in the control group. Figure 1 depicts the six regions affected by the earthquake in a darker shade. The regions in a lighter shade contain all the control counties, where the earthquake was barely felt. In Appendix A, we include a map with the peak ground acceleration at the county level.

4. Evaluating and analyzing the matched sample

4.1 Covariate balance

Before matching, there were 77 exposed counties and 95 control ones. After using cardinality matching to find the largest pair-matched sample of counties that is balanced, we obtained 59 pairs of exposed and control counties. In this section, we describe and evaluate covariate balance in the matched sample. In Section 5.2 below, we discuss the extent to which the results from this sample can be generalized.

Figure 2 shows the absolute standardized differences in means before and after matching for the 18 mean balanced covariates. We observe that all the differences are smaller than 0.1 pooled standard deviations, as specified by the mean balance constraints (4).

Figure 1: Map of the affected regions.

As discussed in Section 2.3, we include the quintiles of the Human Development Index to balance the distributions. Figure 3 below shows the distribution of the HDI quintiles before and after matching. We constrained the mean and the distribution of this covariate in view of previous evidence of a non-linear (U-shaped) relationship with electoral preferences. In panel (a) of the figure, we observe that before matching there are substantial imbalances in the means by quintiles. In panel (b), we observe that after matching these imbalances were substantially reduced.

Figure 4 shows the vote share trajectories for the center-left coalition in the exposed and control groups before and after matching. Before matching, we observe substantial differences in the vote shares, particularly in the elections before the earthquake. After matching, the trajectories are almost indistinguishable between the exposed and control groups. Balancing trajectories allows us to find matched groups that are similar from a dynamic perspective, because they are formed of counties that had similar voting patterns before the earthquake (Haviland et al. 2008).

Ideally, we would have matched exactly for the 18 covariates in our study. However, with 18 covariates, even if these covariates are discretized to be binary and have only two categories, there are 2^{18} or 262,144 possible combinations of types of units, while in our study there are 172 counties (77 exposed and 95 control ones). Furthermore, if the covariates are discretized to have 5 categories, then there are 5^{18} or 3,814,697,265,625 combinations of types of units,

Figure 2: Standardized differences before and after matching.

and by matching exactly we obtain no pairs. In principle, if we were matching exactly for the 18 covariates, we would be granting the same importance to balancing the 18th interaction of the indicators of the tail category of each of the discretized covariates as to balancing the single indicator of the central category of a particular covariate (aiming to balance the first moment of the covariate, which is the first term that appears in a Taylor expansion of the outcome model). While exact matching is the ideal form of covariate balance, there is a curse of dimensionality in exact matching. As discussed previously, flexible forms of covariate balance ranging from exact matching to mean balance can be incorporated within the framework of (1)–(3) to find the largest matched sample that is balanced by design.

4.2 Effect estimates

We estimate the effect of the earthquake on the incumbent vote share using the inferential methods described in Chapter 2 of Rosenbaum (2002). In particular, we use Wilcoxon's signed-rank test statistic to test the sharp null hypothesis of no treatment effect and derive a point estimate by solving Hodges and Lehmann's estimating equation. We obtain a 95% confidence interval by inverting the test. Through this method, we find that the 2010 earthquake increased the incumbent coalition vote share in the 2013 elections by 1.5

Figure 3: Human Development Index quintiles before and after matching.

Figure 4: Vote share trajectories for the center-left coalition before and after matching.

percentage points, with an associated p -value of 0.003 and a 95% confidence interval of [0.006, 0.024]. Table 1 summarizes these results.

As discussed in Section 3.1, several mechanisms could have led to this effect, such as the strong presence of the government in the affected areas, but those mechanisms remain to be tested. Next, in Section 5, we assess the sensitivity of these findings to biases due to unobserved covariates as well as study their generalizability.

Table 1: Effect of the 2010 earthquake on the center-right coalition vote share in 2013.

Hodges-Lehman point estimate	0.015
95% confidence interval	[0.006, 0.024]
p -value	0.003

5. Assessing the sensitivity and generalizability of the findings

5.1 Rosenbaum bounds

In Section 2.1, we explained that in step 1 of cardinality matching we match for global covariate balance, and thus find the largest matched sample that satisfies the covariate balance requirements specified by the investigator. In step 2 of cardinality matching, we rematch the balanced matched sample from step 1 to reduce heterogeneity in the matched pairs. The rematched sample in step 2 is composed of the same treated and control units found in step 1, and the global covariate balance requirements are preserved; however, the covariate distances between rematched units are not greater and typically smaller. As is argued in Rosenbaum (2005) and Zubizarreta et al. (2014), this translates into less sensitivity to hidden bias if the covariates used to compute the distances are strong predictors of the outcome. In this section, we illustrate this procedure and assess the sensitivity of the findings from Section 4.2.

For this, we use the sensitivity analysis model described in Chapter 4 of Rosenbaum (2002). In this model, each matched pair is composed of two units, i and i' , which are identical on their observed covariates, so $\mathbf{x}_i = \mathbf{x}_{i'}$, but perhaps different in terms of an unobserved covariate, so $u_i \neq u_{i'}$. The probabilities of units i and i' of receiving treatment are denoted by π_i and $\pi_{i'}$, respectively. The model posits that the odds of receiving treatment for the two units may differ at most by a factor of $\Gamma \geq 1$ because of the unobserved covariate. More precisely, the model posits that

$$\frac{1}{\Gamma} \leq \frac{\pi_i/(1 - \pi_i)}{\pi_{i'}/(1 - \pi_{i'})} \leq \Gamma \quad (9)$$

for all matched pairs (i, i') . Clearly, if $\Gamma = 1$, then $\pi_i = \pi_{i'}$, and there is no hidden bias in the study. However, if $\Gamma > 1$, then $\pi_i \neq \pi_{i'}$, and for each $\Gamma > 1$, there is a range of possible inferences summarized by two extreme-case p -values. Using this model, we find the largest value of Γ such that the extreme-case p -values reject the null hypothesis of no treatment effect. Naturally, the larger is Γ , the less sensitive are the effect estimates to the influence of an unobserved covariate. In this sense, Γ measures the degree of departure from a study that is free of biases due to unobserved covariates.

As previous research has documented, Chile has low levels of electoral volatility and stable voter choices over time. In other words, within any county, the previous election's vote shares tend to be strong predictors of the vote shares in the next election. For example, the correlation coefficient of the vote share of the incumbent coalition between 2009 and 2013 is 0.85. In step 2 of cardinality matching, we rematch the balanced matched sample from

step 1 to minimize the covariate distance between matched pairs in terms of the incumbent coalition's vote shares in 2009. In step 1, after using cardinality matching for balance, the total sum of covariate distances between matched pairs is 4.48. In step 2, after rematch for homogeneity, this distance decreases to 1.54. The distribution of exposed-minus-control response differences between matched pairs after matching for balance and rematch for homogeneity is shown in Figure 5. We can observe that the exposed-minus-control response differences between matched pairs after rematch for homogeneity are considerably less dispersed.

Figure 5: Distribution of exposed-minus-control response differences between matched pairs after matching for balance (step 1) and rematch for homogeneity (step 2).

Table 2 shows the highest value of the extreme-case p -values for each value of Γ in the matched sample after matching for balance (step 1) and rematch the matched sample for homogeneity (step 2). After matching for balance, the results become sensitive for $\Gamma = 1.00$, whereas after rematch for homogeneity they become sensitive for $\Gamma = 1.50$. In other words, in order for an unobserved covariate to explain away our effect estimate, it would need to increase two identical counties' odds of exposure to the earthquake by a factor of 1.50. This illustrates the value of having pretreatment measures of the outcome in matching, because they tend to be strongly related to the outcome.

Table 2: Upper bounds on the one-sided p -value testing the null hypothesis of no treatment effect using Wilcoxon's signed-rank statistic after steps 1 and 2 of cardinality matching.

Γ	Step 1	Step 2
1.00	0.08	0.002
1.10	0.13	0.004
1.20	0.21	0.01
1.30	0.29	0.02
1.40	0.38	0.03
1.50	0.46	0.05

As discussed in Rosenbaum and Silber (2009), the parameter Γ assumes that the unobserved covariate is a nearly perfect predictor of the outcome, but Γ can be amplified into two parameters Λ and Δ that measure the strength of the relationship between the unobserved covariate and the treatment, and between the unobserved covariate and the outcome, respectively. The results show that $\Gamma = 1.50$ is equivalent to an unobserved covariate that simultaneously increases the odds of exposure to the earthquake and of a greater vote share by $\Lambda = \Delta = 2.62$ times. Similarly, $\Gamma = 1.50$ is equivalent to an unobserved covariate that increases the odds of exposure to the earthquake by $\Lambda = 6$ (or 1.78) times and the odds of a greater vote share by $\Delta = 1.78$ (or 6) times. In other words, for an unobserved covariate to explain away our estimated effect of 1.5%, it would need to have a moderately large effect on the odds of receiving treatment and of exhibiting a positive response.

As an additional robustness check, in Appendix C we conduct a falsification test for an effect that we know to be absent: the effect of the earthquake before it happened. This test provides additional assurance about the effect of the earthquake on electoral outcomes.

5.2 Profile of the matched sample

In order to assess the extent to which the outcome results can be generalized, we present a profile of the matched sample and contrast it to the sample of all the counties in Chile.⁸ As we can see from Table 3 below, in comparison to all the counties in Chile, the matched counties are slightly less populated and more leftist. Thus, some differences between the groups exist, but only in a few covariates, and none of these differences are particularly large.

Table 3: Description of the samples of counties.

Covariate	All	Exposed matched	Control matched
Center-right coalition vote share 1999	0.49	0.47	0.48
Center-right coalition vote share 2005	0.49	0.49	0.48
Center-right coalition vote share 2009	0.44	0.42	0.43
Center-left coalition vote share 1999	0.47	0.49	0.48
Center-left coalition vote share 2005	0.46	0.47	0.48
Center-left coalition vote share 2009	0.33	0.37	0.36
Alternative candidates vote share 1999	0.04	0.04	0.04
Alternative candidates vote share 2005	0.04	0.04	0.04
Alternative candidates vote share 2009	0.23	0.21	0.21
Population 2002	46,540	37,991	34,411
Percentage of rural population 2002	0.38	0.42	0.45
Human development index 2003	0.69	0.68	0.67
Human development index 2003 5th quintile	0.21	0.15	0.14
Human development index 2003 4th quintile	0.19	0.15	0.17
Human development index 2003 3rd quintile	0.20	0.24	0.27
Human development index 2003 2nd quintile	0.20	0.27	0.25
Human development index 2003 1st quintile	0.20	0.19	0.17
Poverty rate 2009	0.17	0.17	0.16

⁸ Chile has 346 counties, but we have excluded a few municipalities that had incomplete data either because they are too small or because they were divided or created after 2004. As a result, we focus on 326 counties.

5.3 Building a representative matched sample

While our balanced matched sample is not very different from the “population” of all counties in Chile, in this section we show how constraints like (7) and (8) can be used to find a matched sample that is not only balanced but also has a similar structure to a target population on observed covariates. Here the target population corresponds to all counties in Chile. In other words, here we estimate the average impact of the earthquake had it hit all the counties in Chile.

By imposing constraints (7) and (8) to the matching problem, we find a matched sample of 50 exposed and control counties that are very similar to all the counties in Chile.⁹ Table 4 compares the means of the exposed and control counties in the matched sample to the means of all the counties in Chile. As required, the absolute standardized differences in means between the matched exposed counties and all the counties, and between the matched control counties and all the counties, are all smaller than 0.05. As a result, the absolute standardized differences in means between the matched exposed and control counties are at most 0.1.

Table 4: Description of the representative matched sample.

Covariate	All	Exposed matched	Control matched
Center-right coalition vote share 1999	0.49	0.49	0.49
Center-right coalition vote share 2005	0.49	0.50	0.49
Center-right coalition vote share 2009	0.44	0.45	0.44
Center-left coalition vote share 1999	0.47	0.47	0.47
Center-left coalition vote share 2005	0.46	0.46	0.46
Center-left coalition vote share 2009	0.33	0.33	0.33
Alternative candidates vote share 1999	0.04	0.04	0.04
Alternative candidates vote share 2005	0.04	0.04	0.04
Alternative candidates vote share 2009	0.23	0.22	0.23
Population 2002	46,540	45,485	47,315
Percentage of rural population 2002	0.38	0.37	0.39
Human development index 2003	0.69	0.69	0.68
Human development index 2003 5th quintile	0.21	0.20	0.20
Human development index 2003 4th quintile	0.19	0.20	0.20
Human development index 2003 3rd quintile	0.20	0.20	0.20
Human development index 2003 2nd quintile	0.20	0.20	0.20
Human development index 2003 1st quintile	0.20	0.20	0.20
Poverty rate 2009	0.17	0.17	0.17

Table 5 displays the impact of the earthquake on the center-right vote share. We follow the same estimation strategy as in Section 4.2. The results obtained using the new matched sample are congruent with the findings from the old matched sample: the earthquake, and subsequent reconstruction process, had a positive and significant impact on the incumbent vote share.

Table 5: Effect of the 2010 earthquake on the center-right coalition vote share in 2013.

Hodges-Lehman point estimate	0.020
95% confidence interval	[0.004, 0.035]
<i>p</i> -value	0.016

⁹These are all the counties in Chile with available data.

6. Comparison to other matching methods

We now compare our results with those obtained through propensity score matching (PSM) and coarsened exact matching (CEM). For PSM, we estimated the propensity score using logistic regression, including the covariates as linear terms in the logistic regression model. With these estimated scores, we conducted two analyses: one using all the available exposed counties, and another dropping those exposed (control) counties without a close counterpart in the control (exposed) group. In both cases we used the package `optmatch` in R (Hansen 2007).

In the first analysis, we obtained 77 matched pairs of counties, but 13 out of the 18 covariates presented substantial imbalances in means (see Figure 6). In the second analysis, we obtained 59 matched pairs of counties, but still 8 out of the 18 covariates exhibited substantial imbalances (Figure 7). Arguably, these imbalances could be improved, but this would require a considerable amount of iteration, as summarized in Algorithm 1.

Figure 6: Standardized differences before and after matching on the estimated propensity score without dropping any counties.

For CEM we used the package `cem` in R. We first used the default coarsening, but obtained no matched pairs. We then relaxed the coarsening to 9, 8, 7, 6, and 5 bins per covariate, but we still obtained no matched pairs. With 4 bins, we obtained 6 matched pairs, and with 3 and 2 bins we obtained 10 and 30 pairs, respectively. In Figure 8 we show balance for the 30-pair match. Though the results are better than those using PSM, 5 out of 18 covariates still have substantial mean imbalances.

By contrast, in our matched sample we obtained 59 matched pairs of counties and all the covariates were balanced not only for their means but in some cases also for their distributions. Ideally we would have matched exactly for every covariate, but this is either

Figure 7: Standardized differences before and after matching on the estimated propensity score after dropping those exposed (control) counties without a counterpart in the control (exposed) group.

an infeasible or very difficult constraint in our case study. In other studies with fewer covariates, exact matching may be possible. This illustrates the flexibility of the covariate balance constraints in cardinality matching, ranging from exact matching to mean balance and including various forms of aggregate distributional balance.

In these analyses, we used R (version 3.4.4) and the running times were as follows: `cem` (version 1.1.19), 0.17 sec; `designmatch` (version 0.3.1), 0.81 sec with `gurobi` (version 8) and 14.43 sec with `Rglpk` (version 0.6-3); `optmatch` (version 0.9-8), 0.31 sec when matching all the treated observations, and 1.06 sec when allowing some of the treated observations to be dropped. See Appendix D for an illustration of cardinality matching using `designmatch` for R.

7. Summary and remarks

In this article, we used cardinality matching to find the largest matched sample that is balanced as designed, before matching, by the investigator. Cardinality matching addresses the problem of limited overlap in covariate distributions using the original covariates, instead of a summary of them, and balances covariates both directly and flexibly. From a statistical standpoint, under a homoskedastic constant additive treatment effect outcome model, cardinality matching minimizes the variance of a difference-in-means effect estimator that is approximately unbiased according to the covariate balance constraints of the matching problem. From a computational standpoint, cardinality matching takes advantage of recent

Figure 8: Standardized differences before and after CEM.

developments in mathematical programming to solve relatively large matching problems either exactly or approximately. Finally, as illustrated in Section 5, after finding the largest matched sample that is balanced by design, the matched sample can be rematched to reduce heterogeneity in matched groups and thus reduce sensitivity to unobserved covariates. Cardinality matching, as well as other matching methods for instrumental variables, discontinuity designs, and design of experiments, is implemented in the new statistical package `designmatch` for R (Zubizarreta et al. 2017).

We applied this method to estimate the effects of the 2010 earthquake in Chile in electoral outcomes. We provide evidence that voters rewarded the incumbent candidate after the 2010 earthquake in Chile. Contrary to what is expected from the theory by Achen and Bartels (2016), disaster victims do not necessarily punish the incumbent for events that deteriorate their living conditions. Our findings are in line with the theory by Healy and Malhotra (2010) that states that citizens can either punish or reward incumbents based on their performance handling the consequences of a natural catastrophe.

Acknowledgments

We thank Jake Bowers, Zach Branson, Shigeo Hirano, Gary King, Christopher Lucas, and Luke Miratrix for useful comments and suggestions. We also thank the support from a grant from the Alfred P. Sloan Foundation. All remaining errors are our responsibility.

Appendix A. Peak ground acceleration

The following map shows the peak ground acceleration at the county level. A darker shade represents a higher intensity.

Figure 9: Peak ground acceleration at the county level.

Appendix B. Missing data

There are a few counties that do not have data for some covariates because they are too small, or they were divided or created after 2004. All these counties are excluded from the matching and posterior analysis of outcomes. These are: Antartica, Laguna Blanca,

Tortel, Guaitecas, Lago Verde, Timaukel, Torres del Paine, Río Verde, O'Higgins, Chaitén, Juan Fernandez, Isla de Pascua, Alto Hospicio, Hualpén, Alto Bío Bío, Cholchol, Iquique, Talcahuano, Santa Bárbara, and Nueva Imperial.

Appendix C. A falsification test

As an additional robustness check, we conduct a falsification test for an effect that we know to be absent. For example, we know the 2010 earthquake should have no effect on elections before 2010. Furthermore, we exploit the fact that the same center-left candidate (Michelle Bachelet) ran for president in 2005 (pre-earthquake) and 2013 (post-earthquake). This creates an opportunity for a more credible before/after comparison because we can control for the candidate's personal characteristics.

We conduct the falsification test in two stages. In stage 1, we test the effect of the earthquake on the vote share of the (challenger) center-left candidate in 2013. We use the same matched sample as before, but now rematching the units with the vote share of the center-left wing candidates in previous elections. Because we already know that the incumbent improved her vote share, we would expect the earthquake to have weakened the challenger's electoral performance.

In stage 2, we use the 59 matched treated counties obtained in the matched sample to conduct a falsification test for the 2005 election. The falsification control group is optimally defined by selecting the counties that meet the balance constraints for the pre-2005 covariates. The same 59 counties are used in the matched exposed group for the 2013 and 2005 elections to produce more credible comparisons, but the control group counties can be different since they are selected to achieve balance. Because we know that there was no 2005 earthquake, we would expect to see no differences between the matched exposed and control groups.

As expected, the results from stage 1 show that the earthquake decreases the vote share of the (challenger) center-left candidate in the 2013 election, and stage 2 shows that there are no differences between the matched exposed and control groups in the 2005 election. This falsification analysis gives us more confidence about the electoral effect of the 2010 earthquake in Chile: if there were hidden biases we would have observed a difference between both groups of counties before the catastrophe.

Table 6: Wilcoxon signed-rank test: Center-left candidate vote share in 2005 and 2013.

	2005 Election	2013 Election
Hodges-Lehman point estimate	-0.013	-0.030
95% confidence interval	[-0.038, 0.014]	[-0.046, -0.013]
<i>p</i> -value	0.330	0.001

Appendix D. Cardinality matching with `designmatch` for R

In this appendix, we show how to use cardinality matching with the `designmatch` package for R. First, we use `cardmatch` to find the largest pair-matched sample that is balanced. Second, we use `distmatch` to rematch the balanced sample found by `cardmatch` to minimize the total sum of covariate distances, while preserving balance on aggregate. Finally, we show how to use `cardmatch` to find the largest pair-matched sample that is both balanced and representative of a population distribution of particular interest.

D.1 Finding the largest pair-matched sample that is balanced

First, we load the packages `designmatch` and `gurobi`. `gurobi` greatly enhances the performance of `designmatch`, and is specially useful in large data sets, but it requires special installation. For details about the installation, please read the first sections of https://cran.r-project.org/web/packages/prioritizr/vignettes/gurobi_installation.html.

```
> library(designmatch)
> library(gurobi)
```

Now, we load, sort, and attach the data set `d` in `earthquake.RData`.

```
> load("earthquake.RData")
> d = d[order(d$t_ind, decreasing=TRUE), ]
> attach(d)
```

The exposure (treatment) indicator is given by `t_ind`. There are 77 and 99 exposed and unexposed counties, respectively.

```
> table(t_ind)
t_ind
 0  1
95 77
```

Next, we define the moment balance requirements; in this case, to balance the means of the covariates up to 0.1 standardized differences. The covariates are the vote shares in the three previous presidential elections of the center-right coalition (`lavin99_p`, `lavinpinera05_p`, `pinera09_p`), the center-left coalition (`lagos99_p`, `bagelet05_p`, `frei09_p`), and the alternative candidates (`nontraditional99_p`, `nontraditional05_p`, `nontraditional09_p`); the population and proportion of rural population in the previous census (`pop2002`, `rural02_p`); the human development index (`IDH03`) plus indicators for its quantiles (`IDH5`, `IDH4`, `IDH3`, `IDH2`, `IDH1`) in 2003; and the poverty rate in 2009 (`poverty09`).

```
> mom_covs = cbind(lavin99_p, lavinpinera05_p, pinera09_p,
+ lagos99_p, bagelet05_p, frei09_p,
+ nontraditional99_p, nontraditional05_p, nontraditional09_p,
+ pop2002, rural02_p,
+ IDH03, IDH5, IDH4, IDH3, IDH2, IDH1,
+ poverty09)
```

```
> mom_tols = absstdif(mom_covs, t_ind, .1)
> mom = list(covs = mom_covs, tols = mom_tols, targets = NULL)
```

Finally, we specify the solver parameters.

```
> t_max = 60*30
> name = "gurobi"
> approximate = 0
> solver = list(name = name, t_max = t_max, approximate = approximate,
+ round_cplex = 0, trace = 1)
```

With this, we may use `cardmatch` to find largest pair-matched sample for which the means of the covariates are balanced up to 0.1 standardized differences.

```
> out_1 = cardmatch(t_ind, mom = mom, solver = solver)
Building the matching problem...
Gurobi optimizer is open...
Finding the optimal matches...
Optimal matches found
```

We extract the indices of the matched treated and control units, and count the number of units that were matched.

```
> t_id_1 = out_1$t_id
> c_id_1 = out_1$c_id
> length(t_id_1)
[1] 59
> length(c_id_1)
[1] 59
```

We assess balance by checking the standardized differences in means after matching. As we can see from the second to last column, after matching all the means are at most 0.1 pooled standard deviations apart.

```
> meantab(mom_covs, t_ind, t_id_1, c_id_1, digits=2)
```

	Mis	Min	Max	Mean T	Mean C	Std Dif	P-val
lavin99_p	0	0.21	0.86	0.47	0.48	-0.10	0.60
lavinpinera05_p	0	0.27	0.87	0.49	0.48	0.10	0.58
pinera09_p	0	0.24	0.78	0.42	0.43	-0.09	0.61
lagos99_p	0	0.12	0.74	0.49	0.48	0.10	0.60
bachalet05_p	0	0.11	0.68	0.47	0.48	-0.09	0.62
frei09_p	0	0.07	0.55	0.37	0.36	0.10	0.57
nontraditional99_p	0	0.02	0.12	0.04	0.04	0.05	0.77
nontraditional05_p	0	0.02	0.11	0.04	0.04	-0.10	0.59
nontraditional09_p	0	0.12	0.37	0.21	0.21	-0.01	0.97
pop2002	0	1204.00	220826.00	37991.29	34411.47	0.07	0.70
rural02_p	0	0.00	1.00	0.42	0.45	-0.10	0.58
IDH03	0	0.51	0.79	0.68	0.67	0.01	0.97
IDH5	0	0.00	1.00	0.15	0.14	0.05	0.80

IDH4	0	0.00	1.00	0.15	0.17	-0.05	0.80
IDH3	0	0.00	1.00	0.24	0.27	-0.08	0.68
IDH2	0	0.00	1.00	0.27	0.25	0.04	0.84
IDH1	0	0.00	1.00	0.19	0.17	0.04	0.81
poverty09	0	0.03	0.37	0.17	0.16	0.09	0.59

D.2 Rematching the largest matched sample that is balanced

Now that we have found the largest pair-matched sample that is mean balanced, we will rematch its matched pairs to minimize heterogeneity. For this we will use the function `distmatch`, but first we need a new treatment indicator and a distance matrix for the units matched by `cardmatch`.

```
> t_ind_2 = t_ind[c(t_id_1, c_id_1)]
> table(t_ind_2)
t_ind_2
 0  1
59 59
```

To build the distance matrix, we use strong predictors of the outcome (here, the electoral results before the earthquake).

```
> dist_mat_2 = abs(outer(pinera09_p[t_id_1], pinera09_p[c_id_1], "-"))
```

Using the previous solver options, we use `distmatch` to find the match that minimizes the total sum of distances as given by `dist_mat_2` between the pairs found by `cardmatch`.

```
> out_2 = distmatch(t_ind_2, dist_mat_2, solver)
Building the matching problem...
Gurobi optimizer is open...
Finding the optimal matches...
Optimal matches found
```

We generate the new indices of the treated units and matched controls.

```
> t_id_2 = t_id_1[out_2$t_id]
> c_id_2 = c_id_1[out_2$c_id-length(out_2$c_id)]
```

After rematching, we check that the total sum of covariate distances has been reduced, but that mean balance has been preserved.

```
> distances_step_1 = sum(diag(dist_mat_2))
> distances_step_2 = sum(diag(dist_mat_2[out_2$t_id, out_2$c_id-length(out_2$c_id)]))
> distances_step_1
[1] 4.480845
> distances_step_2
[1] 1.539447

> meantab(mom_covs, t_ind, t_id_2, c_id_2, digits=2)
```

	Mis	Min	Max	Mean T	Mean C	Std Dif	P-val
lavin99_p	0	0.21	0.86	0.47	0.48	-0.10	0.60
lavinpinera05_p	0	0.27	0.87	0.49	0.48	0.10	0.58
pinera09_p	0	0.24	0.78	0.42	0.43	-0.09	0.61
lagos99_p	0	0.12	0.74	0.49	0.48	0.10	0.60
batchelet05_p	0	0.11	0.68	0.47	0.48	-0.09	0.62
frei09_p	0	0.07	0.55	0.37	0.36	0.10	0.57
nontraditional99_p	0	0.02	0.12	0.04	0.04	0.05	0.77
nontraditional05_p	0	0.02	0.11	0.04	0.04	-0.10	0.59
nontraditional09_p	0	0.12	0.37	0.21	0.21	-0.01	0.97
pop2002	0	1204.00	220826.00	37991.29	34411.47	0.07	0.70
rural02_p	0	0.00	1.00	0.42	0.45	-0.10	0.58
IDH03	0	0.51	0.79	0.68	0.67	0.01	0.97
IDH5	0	0.00	1.00	0.15	0.14	0.05	0.80
IDH4	0	0.00	1.00	0.15	0.17	-0.05	0.80
IDH3	0	0.00	1.00	0.24	0.27	-0.08	0.68
IDH2	0	0.00	1.00	0.27	0.25	0.04	0.84
IDH1	0	0.00	1.00	0.19	0.17	0.04	0.81
poverty09	0	0.03	0.37	0.17	0.16	0.09	0.59

Finally, we save the matched and rematched samples.

```
> d_match = d[c(t_id_1, c_id_1), ]
> d_rematch = d[c(t_id_2, c_id_2), ]
```

To analyze the outcomes, we may use Wilcoxon's signed rank test.

```
> library(exactRankTests)

> wilcox.exact(d_rematch$matthei13_p[d_rematch$t_ind==1],
+ d_rematch$matthei13_p[d_rematch$t_ind==0], paired=TRUE, conf.int=TRUE)
```

Asymptotic Wilcoxon signed rank test

```
data: d_rematch$matthei13_p[d_rematch$t_ind == 1] and
d_rematch$matthei13_p[d_rematch$t_ind == 0]
V = 1277, p-value = 0.003088
alternative hypothesis: true mu is not equal to 0
95 percent confidence interval:
0.005931411 0.024281858
sample estimates:
(pseudo)median
0.01536187
```

D.3 Building a balanced and representative matched sample

Following sections 2.4 and 5.3, we can build a representative matched sample with the argument `mom_targets` in `mom`. `mom_targets` defines the moments of the target distribution. In this example, the target distribution is the one of all the counties in Chile, and `mom_targets` are the first moments of this distribution. In principle, one can target higher-order univariate and multivariate moments by appropriately transforming the covariates (see Zubizarreta 2012 for a discussion). We use `cardmatch` as follows.

```
> mom_covs = cbind(lavin99_p, lavinpinera05_p, pinera09_p,
+ lagos99_p, bachelet05_p, frei09_p,
+ nontraditional99_p, nontraditional05_p, nontraditional09_p,
+ pop2002, rural02_p,
+ IDH03, IDH5, IDH4, IDH3, IDH2, IDH1,
+ poverty09)
> mom_tols = absstdif(mom_covs, t_ind, .05)
> mom_targets = c(0.49443, 0.49448, 0.44193, 0.46591, 0.46266, 0.33193, 0.03965, 0.04285,
+ 0.22614, 46540.48773, 0.38017, 0.68711, 0.21166, 0.19018, 0.19939, 0.19939,
+ 0.19939, 0.16805)
> mom = list(covs = mom_covs, tols = mom_tols, targets = mom_targets)
> out_3 = cardmatch(t_ind, mom = mom, solver = solver)
Building the matching problem...
Gurobi optimizer is open...
Finding the optimal matches...
Optimal matches found
```

Like before, we extract the indices of the matched treated and control units, and count the number of matched pairs.

```
> t_id_3 = out_3$t_id
> c_id_3 = out_3$c_id
> length(t_id_3)
[1] 25
> length(c_id_3)
[1] 25
```

We check balance, and observe that the matched groups are not only balanced but similar to the target distribution.

```
> cbind(meantab(mom_covs, t_ind, t_id_3, c_id_3, digits=3)[, 4:5], round(mom_targets, 3))
```

	Mean T	Mean C	mom_targets
lavin99_p	0.490	0.491	0.494
lavinpinera05_p	0.499	0.494	0.494
pinera09_p	0.445	0.442	0.442
lagos99_p	0.470	0.470	0.466
bachelet05_p	0.459	0.462	0.463
frei09_p	0.331	0.332	0.332

nontraditional99_p	0.040	0.039	0.040
nontraditional05_p	0.042	0.043	0.043
nontraditional09_p	0.223	0.226	0.226
pop2002	45485.120	47314.960	46540.488
rural02_p	0.372	0.395	0.380
IDH03	0.685	0.685	0.687
IDH5	0.200	0.200	0.212
IDH4	0.200	0.200	0.190
IDH3	0.200	0.200	0.199
IDH2	0.200	0.200	0.199
IDH1	0.200	0.200	0.199
poverty09	0.166	0.167	0.168

In a similar way to section D.2, one may proceed by rematching this balanced sample for homogeneity.

References

- Abadie, A., Drukker, D., Herr, J. L., Imbens, G. W., et al. (2004). Implementing matching estimators for average treatment effects in stata. *Stata Journal*, 4:290–311.
- Achen, C. H. and Bartels, L. (2004). Blind retrospection: Electoral responses to drought, flu, and shark attacks. Working Paper Center for Advanced Study in the Social Sciences. Available from: goo.gl/W5yWLv.
- Achen, C. H. and Bartels, L. M. (2016). *Democracy for realists: Why elections do not produce responsive government*. Princeton University Press.
- Altman, D. (2004). Redrawing the chilean electoral map: The presence of socioeconomic and gender factors in the ballot box. *Revista de Ciencia Politica*, 24(2):49.
- Arana, I. (2016). Aftershocks of pinochet’s constitution: The chilean postearthquake reconstruction. *Latin American Perspectives*, page 0094582X16637145.
- Armitage, P. (1982). The role of randomization in clinical trials. *Statistics in Medicine*, 1(4):345–352.
- Aronow, P. M. and Samii, C. (2016). Does regression produce representative estimates of causal effects? *American Journal of Political Science*, 60(1):250–267.
- Austin, P. C. (2011). An introduction to propensity score methods for reducing the effects of confounding in observational studies. *Multivariate behavioral research*, 46(3):399–424.
- Bechtel, M. M. and Hainmueller, J. (2011). How lasting is voter gratitude? an analysis of the short-and long-term electoral returns to beneficial policy. *American Journal of Political Science*, 55(4):852–868.
- Bennett, M., Vielma, J. P., and Zubizarreta, J. R. (2018). Building representative matched samples with multi-valued treatments. *Working Paper*.

- Calvo, E. and Murillo, M. V. (2012). When parties meet voters assessing political linkages through partisan networks and distributive expectations in argentina and chile. *Comparative Political Studies*, 46(7):851–882.
- Carlin, R. E., Love, G. J., and Zechmeister, E. J. (2014). Natural disaster and democratic legitimacy the public opinion consequences of chile’s 2010 earthquake and tsunami. *Political Research Quarterly*, 67(1):3–15.
- Chen, J. (2013). Voter partisanship and the effect of distributive spending on political participation. *American Journal of Political Science*, 57(1):200–217.
- Choi, C. Chile quake and tsunami dramatically altered ecosystems [online]. (2012). Available from: goo.gl/2pAc1o.
- Cochran, W. and Rubin, D. (1973). Controlling bias in observational studies: A review. *Sankhyā: The Indian Journal of Statistics, Series A*, pages 417–446.
- Cochran, W. G. and Chambers, S. P. (1965). The planning of observational studies of human populations. *Journal of the Royal Statistical Society: Series A*, 128(2):234–266.
- Crump, R. K., Hotz, V. J., Imbens, G. W., and Mitnik, O. A. (2009). Dealing with limited overlap in estimation of average treatment effects. *Biometrika*, 96(1):187–199.
- Diamond, A. and Sekhon, J. S. (2013). Genetic matching for estimating causal effects: A general multivariate matching method for achieving balance in observational studies. *Review of Economics and Statistics*, 95(3):932–945.
- Dorn, H. F. (1953). Philosophy of inferences from retrospective studies. *American Journal of Public Health and the Nations Health*, 43(6):677–683.
- Fogarty, C. B., Mikkelsen, M. E., Gaieski, D. F., and Small, D. S. (2016). Discrete optimization for interpretable study populations and randomization inference in an observational study of severe sepsis mortality. *Journal of the American Statistical Association*, 111(514):447–458.
- Gasper, J. T. and Reeves, A. (2011). Make it rain? retrospection and the attentive electorate in the context of natural disasters. *American Journal of Political Science*, 55(2):340–355.
- González, J. I. (1999). Geografía electoral de chile: Comportamiento del electorado chileno entre 1932-1992. *Estudios Geográficos*, 60(234):121–138.
- Government of Chile. Plan de reconstrucción: Terremoto y maremoto del 27 de febrero 2010 [online]. (2010). Available from: goo.gl/iP0FcM.
- Hainmueller, J. (2012). Entropy balancing for causal effects: a multivariate reweighting method to produce balanced samples in observational studies. *Political Analysis*, 20(1):25–46.
- Hansen, B. (2007). Flexible, optimal matching for observational studies. *R News*, 7:18–24.

- Haviland, A., Nagin, D., and Rosenbaum, P. (2007). Combining propensity score matching and group-based trajectory analysis in an observational study. *Psychological Methods*, 12(3):247.
- Haviland, A., Nagin, D. S., Rosenbaum, P. R., and Tremblay, R. E. (2008). Combining group-based trajectory modeling and propensity score matching for causal inferences in nonexperimental longitudinal data. *Developmental psychology*, 44(2):422.
- Healy, A. and Malhotra, N. (2010). Random events, economic losses, and retrospective voting: Implications for democratic competence. *International Quarterly Journal of Political Science*, 5(2):193–208.
- Hinrichs, R., Jones, L., and Stanley, E. Report on the 2010 chilean earthquake and tsunami response [online]. (2011). Available from: <http://pubs.usgs.gov/of/2011/1053/of2011-1053.pdf>.
- Hirshberg, D. A. and Zubizarreta, J. R. (2018). Consistent, representative, and scalable matching for large observational studies. *Working Paper*.
- Ho, D. E., Imai, K., King, G., and Stuart, E. A. (2007). Matching as nonparametric preprocessing for reducing model dependence in parametric causal inference. *Political Analysis*, 15(3):199–236.
- Hsu, J. Y., Zubizarreta, J. R., Small, D. S., and Rosenbaum, P. R. (2015). Strong control of the familywise error rate in observational studies that discover effect modification by exploratory methods. *Biometrika*, 102(4):767–782.
- Iacus, S. M., King, G., and Porro, G. (2012). Causal inference without balance checking: Coarsened exact matching. *Political analysis*, 20(1):1–24.
- Imai, K. and Ratkovic, M. (2014). Covariate balancing propensity score. *Journal of the Royal Statistical Society: Series B*, 76(1):243–263.
- Imbens, G. W. (2015). Matching methods in practice: Three examples. *Journal of Human Resources*, 50(2):373–419.
- Kilcioglu, C., Zubizarreta, J. R., et al. (2016). Maximizing the information content of a balanced matched sample in a study of the economic performance of green buildings. *The Annals of Applied Statistics*, 10(4):1997–2020.
- King, G., Lucas, C., and Nielsen, R. (2016). The balance-sample size frontier in matching methods for causal inference. *American Journal of Political Science*.
- LaLonde, R. J. (1986). Evaluating the econometric evaluations of training programs with experimental data. *The American Economic Review*, pages 604–620.
- Lazarev, E., Sobolev, A., Soboleva, I. V., and Sokolov, B. (2014). Trial by fire: A natural disaster’s impact on support for the authorities in rural russia. *World Politics*, 66(04):641–668.

- Little, R. J. and Rubin, D. B. (2014). *Statistical Analysis with Missing Data*. John Wiley & Sons.
- López, M. Á. (2004). Conducta electoral y estratos económicos: El voto de los sectores populares en Chile. *Política*, 43:285–298.
- Luna, J. P. (2010). Segmented party–voter linkages in Latin America: The case of the UDI. *Journal of Latin American Studies*, 42(02):325–356.
- Maldonado, L., Kronmüller, E., and Gutierrez, I. (2016). Estrategia para la inferencia causal y planificación de estudios observacionales en las ciencias sociales: el caso de Chile post erupción del 2008. *Revista de ciencia política (Santiago)*, 36(3):797–827.
- McClean, D. Chile still living with quake effects [online]. (2012). Available from: <http://www.unisdr.org/archive/25366>.
- Navia, P., Izquierdo, J. M., and Morales, M. (2008). Voto cruzado en Chile: ¿por qué Bachelet obtuvo menos votos que la Concertación en 2005? *Política y Gobierno*, 15(1):35–73.
- Roberts, K. M. (2013). Market reform, programmatic (de)alignment, and party system stability in Latin America. *Comparative Political Studies*, 46(11):1422–1452.
- Rosenbaum, P. R. (1989). Optimal matching for observational studies. *Journal of the American Statistical Association*, 84(408):1024–1032.
- Rosenbaum, P. R. (2002). *Observational Studies*. Springer.
- Rosenbaum, P. R. (2004). Design sensitivity in observational studies. *Biometrika*, 91(1):153–164.
- Rosenbaum, P. R. (2005). Heterogeneity and causality: Unit heterogeneity and design sensitivity in observational studies. *The American Statistician*, 59(2):147–152.
- Rosenbaum, P. R. (2010). *Design of Observational Studies*. Springer.
- Rosenbaum, P. R. (2012). Optimal matching of an optimally chosen subset in observational studies. *Journal of Computational and Graphical Statistics*, 21(1):57–71.
- Rosenbaum, P. R. (2015). Multivariate matching methods. In S. Kotz, C. R. R. and Banks, D., editors, *Encyclopedia of Statistical Sciences*, volume 2, pages 435–438. John Wiley.
- Rosenbaum, P. R., Ross, R. N., and Silber, J. H. (2007). Minimum distance matched sampling with fine balance in an observational study of treatment for ovarian cancer. *Journal of the American Statistical Association*, 102(477):75–83.
- Rosenbaum, P. R. and Rubin, D. B. (1983). The central role of the propensity score in observational studies for causal effects. *Biometrika*, 70(1):41–55.
- Rosenbaum, P. R. and Rubin, D. B. (1985). Constructing a control group using multivariate matched sampling methods that incorporate the propensity score. *The American Statistician*, 39(1):33–38.

- Rosenbaum, P. R. and Silber, J. H. (2009). Amplification of sensitivity analysis in observational studies. *Journal of the American Statistical Association*, 104(488):1398–1405.
- Rubin, D. B. (1973). Matching to remove bias in observational studies. *Biometrics*, pages 159–183.
- Rubin, D. B. (2006). *Matched Sampling for Causal Effects*. Cambridge University Press.
- Sinclair, B., Hall, T. E., and Alvarez, R. M. (2011). Flooding the vote: Hurricane katrina and voter participation in new orleans. *American politics research*, 39(5):921–957.
- Stuart, E. A. (2010). Matching methods for causal inference: A review and a look forward. *Statistical Science*, 25(1):1–21.
- Visconti, G. (2018). *Political Preferences in Adverse Conditions*. PhD thesis, Columbia University.
- Zubizarreta, J. R. (2012). Using mixed integer programming for matching in an observational study of kidney failure after surgery. *Journal of the American Statistical Association*, 107(500):1360–1371.
- Zubizarreta, J. R., Cerdá, M., and Rosenbaum, P. R. (2013). Effect of the 2010 chilean earthquake on posttraumatic stress: Reducing sensitivity to unmeasured bias through study design. *Epidemiology*, 24(1):79–87.
- Zubizarreta, J. R., Kilcioglu, C., and Vielma, J. P. (2017). designmatch: Matched samples that are balanced and representative by design. *R package version 0.3*, 1.
- Zubizarreta, J. R., Paredes, R. D., and Rosenbaum, P. R. (2014). Matching for balance, pairing for heterogeneity in an observational study of the effectiveness of for-profit and not-for-profit high schools in chile. *Annals of Applied Statistics*, 8(1):204–231.