

An Algorithm for Optimal Tapered Matching, With Application to Disparities in Survival

Shoshana R. DANIEL, Katrina ARMSTRONG,
Jeffrey H. SILBER, and Paul R. ROSENBAUM

In a tapered matched comparison, one group of individuals, called the focal group, is compared to two or more nonoverlapping matched comparison groups constructed from one population in such a way that successive comparison groups increasingly resemble the focal group. An optimally tapered matching solves two problems simultaneously: it optimally divides the single comparison population into nonoverlapping comparison groups and optimally pairs members of the focal group with members of each comparison group. We show how to use the optimal assignment algorithm in a new way to solve the optimally tapered matching problem, with implementation in R. This issue often arises in studies of groups defined by race, gender, or other categorizations such that equitable public policy might require an understanding of the mechanisms that produce disparate outcomes, where certain specific mechanisms would be judged illegitimate, necessitating reform. In particular, we use data from Medicare and the SEER Program of the National Cancer Institute as part of an ongoing study of black-white disparities in survival among women with endometrial cancer.

Key Words: Assignment algorithm; Combinatorial optimization; Matched sampling.

1. INTRODUCTION AND REVIEW

1.1 INTRODUCTION: WHAT IS TAPERED MATCHING?

In a tapered matched comparison, one group of individuals, called the focal group, is compared to two or more nonoverlapping matched comparison groups constructed from one population in such a way that successive comparison groups increasingly resemble

Shoshana R. Daniel is Biostatistician, Covance Periapproval Services, Inc., 555 North Lane, Suite 6000, Conshohocken, PA 19428. Katrina Armstrong is Associate Professor of Medicine, University of Pennsylvania School of Medicine, 1233 Blockley Hall, Philadelphia, PA 19104-6021. Jeffrey H. Silber is Professor, The University of Pennsylvania, and Director, Center for Outcomes Research, The Children's Hospital of Philadelphia, 3535 Market Street, Suite 1029, Philadelphia, PA 19104. Paul R. Rosenbaum is Robert G. Putzel Professor, Department of Statistics, The Wharton School, University of Pennsylvania, 473 Jon M. Huntsman Hall, 3730 Walnut Street, Philadelphia, PA 19104-6340 (E-mail for correspondence: rosenbaum@stat.wharton.upenn.edu).

© 2008 American Statistical Association, Institute of Mathematical Statistics,
and Interface Foundation of North America

Journal of Computational and Graphical Statistics, Volume 17, Number 4, Pages 914–924

DOI: 10.1198/106186008X385806

the focal group. “Tapering” refers to this narrowing or successive removal of naturally occurring discrepancies between the focal group and the comparison population. The focal group never changes; it represents the focal group as it actually exists in the population. The successive comparison groups resemble their own population less and less, and resemble the focal population more and more. Tapered matching is used in an effort to understand how and why outcomes in the focal population differ from those in the comparison population, by successively removing some discrepancies that may be responsible for part of the difference in outcomes. Because the comparison groups are nonoverlapping, they may be compared to one another using simple statistical methods. Tapered comparisons of this sort shed some light on questions of the form: When certain baseline discrepancies are removed, how much of the disparity in outcomes is also removed?

An optimally tapered matching solves two problems simultaneously: it optimally divides the single population into $C \geq 2$ nonoverlapping comparison groups and optimally pairs members of the focal group with members of each comparison group. More generally, each member of the focal group may be matched to $k_c \geq 1$ controls in comparison group c , $c = 1, \dots, C$. For discussion of matching with one comparison group and multiple controls from that group, see Smith (1997).

This problem occurs in various contexts, including studies of disparities in outcomes in which one is seeking insight into mechanisms that produce the disparities. In the example, black women with endometrial cancer (the focal group) survive for a shorter period following diagnosis than do white women, and we compare black women to white women who were diagnosed with similar cancer (a wider comparison group) and white women who were similar in diagnosis and also similar in surgical treatment (a narrower comparison group). Because the two white groups do not overlap, they may be compared to each other using conventional statistical methods to ask: If whites who were diagnosed in the same health as blacks also received the same surgery as blacks, could the difference in surgery help in understanding the difference in survival? Unaided by matching, a model fitted to the entire population would give disproportionate weight to the large and comparatively healthy white population, and that model may substantially misrepresent the situation in the small, sicker black population; see Dehejia and Wahba (1999) for discussion of the failure of models unaided by matching. In contrast, it is often useful to apply model-based adjustments to matched samples (Rubin 1979).

We use the assignment algorithm in a new way to solve the optimally tapered matching problem. The assignment algorithm is reviewed in Section 1.2. An example is discussed in Section 2. Advantages of tapered matching in the context of the example are discussed in Section 2.2. The method for optimal tapered matching is discussed in Section 3. Technical details of the example in Section 2 are discussed in Section 3.3. The choice of distance function in tapered matching is discussed in Section 3.4.

1.2 REVIEW: THE OPTIMAL ASSIGNMENT ALGORITHM

The assignment algorithm is one of the oldest combinatorial optimization algorithms, and it solves the following problem. There is an $A \times B$ matrix Δ of nonnegative, possibly infinite “distances,” $\delta_{ab} \geq 0$, $a = 1, \dots, A$, $b = 1, \dots, B$, with $A \leq B$. These

“distances” need not satisfy the triangle inequality, so they need not be distances in the sense of a metric space. One familiar distance is the Mahalanobis distance (Rubin 1979). An *assignment*, $\beta(\cdot)$, of columns to rows pairs each row $a \in \{1, \dots, A\}$ to a different column $\beta(a) \in \{1, \dots, B\}$ with $\beta(a) \neq \beta(a')$ if $a \neq a'$. Let \mathcal{B} be the set containing all $B!/ (B - A)!$ assignments $\beta(\cdot)$. The *total distance*, $\delta\{\beta(\cdot)\}$, for an assignment $\beta(\cdot)$ is sum of the distances for its A row-column pairs, $\delta\{\beta(\cdot)\} = \sum_{a=1}^A \delta_{a,\beta(a)}$. The *optimal assignment problem* is to find any assignment, $\tilde{\beta}(\cdot) \in \mathcal{B}$, that is optimal in the sense of minimizing the total distance over all assignments:

$$\delta\{\tilde{\beta}(\cdot)\} = \min_{\beta(\cdot) \in \mathcal{B}} \delta\{\beta(\cdot)\} \text{ or equivalently } \sum_{a=1}^A \delta_{a,\tilde{\beta}(a)} = \min_{\beta(\cdot) \in \mathcal{B}} \sum_{a=1}^A \delta_{a,\beta(a)}. \quad (1.1)$$

The problem is not trivial: two or more rows, a and a' , may want the same column b —that is, $\delta_{ab} = \min_j \delta_{aj}$ and $\delta_{a'b} = \min_j \delta_{a'j}$ —so the A assignments of columns to rows are interdependent. The assignment algorithm and related procedures have been used for optimal matching in observational studies; see Rosenbaum (1989, 1991), Silber et al. (2001, 2005, 2007), Hansen (2004, 2007), Hansen and Klopfer (2006), Lu (2005), Augurzky and Kluge (2007), and Rosenbaum, Silber, and Ross (2007). One of the fastest algorithms holds an auction, in which two rows a and a' that both want the same column b offer competing bids for it, with the bids adjusting as the auction progresses; see Bertsekas (1981). The `pairmatch` function in Hansen’s (2004, 2007) `optmatch` package in R makes Bertsekas’ very fast Fortran code available from within R. Dell’Amico and Toth (2000) reviewed and compared algorithms to solve the assignment problem. As discussed by Papadimitriou and Steiglitz (1982), polynomial time algorithms for the assignment problem are available, including Kuhn’s Hungarian algorithm, with time bound of $O(B^3)$. Carpaneto and Toth (1980) provided Fortran code implementing the Hungarian method. In SAS, `PROC ASSIGN` solves the assignment problem; see Bergstralh, Kosanke, and Jacobsen (1996).

As we demonstrate in Section 3, for a suitably defined matrix Δ , the problem of optimal tapered matching can be reduced to the assignment problem. Specifically, the assignment algorithm will both partition the comparison population into nonoverlapping comparison groups matched in varying degrees to the focal group and also match each member of the focal group to k_c similar members of comparison group c , $c = 1, \dots, C$. To make the discussion tangible, before discussing the procedure, its application is illustrated in Section 2 using the example that motivated our work.

2. EXAMPLE: DISPARITIES IN SURVIVAL WITH ENDOMETRIAL CANCER

2.1 MATCHING AT DIAGNOSIS AND AT TREATMENT: AN APPLICATION OF TA- PERED MATCHING

In the U.S. Medicare population, black women with endometrial cancer survive for a shorter time following diagnosis than do white women. Why? Biological differences are not inconceivable, but there is naturally concern that the health care system might be

providing inferior care to black women. In principle, all women in the Medicare population have equal access to medical care, but gaps between principle and practice may occur. We examine this using merged data from Medicare and from the Surveillance, Epidemiology, and End Results (SEER) program of the National Cancer Institute, which provides clinical stage, grade and histology.

We begin with the population of black women diagnosed with endometrial cancer, specifically 806 black women diagnosed between 1991 and 2000. We compare them to two optimally matched, nonoverlapping comparison groups: (i) white women who appeared to have similar cancer at diagnosis, and (ii) white women who appeared to have similar cancer at diagnosis and who received similar surgical treatment. The construction of the matched comparison is discussed in general in Section 3 with specifics for this example in Section 3.3. We also compare with the unmatched population of white women with endometrial cancer. This offers several perspectives and answers several questions. Are blacks and whites diagnosed with similarly advanced cancer? (They aren't.) Do blacks and whites with similar cancer at diagnosis have similar survival? (They don't.) Do blacks and whites with similar cancer receive the same surgery? (They don't.) Do blacks and whites with similar cancer who receive the same surgery exhibit similar survival? (They don't.) Do whites who receive the same inferior treatment as blacks have poorer survival than other whites with similar cancer at diagnosis? (No, or at least there is no indication that they do.) The analyses presented here were a pilot study demonstrating the feasibility of the methodology. A clinical analysis along the same lines is currently underway; it will differ in a number of particulars, including greater attention to the extent and nature of comorbid conditions. A reader interested in endometrial cancer rather than in multivariate matching should refer to our clinical analysis.

Table 1 describes the four groups of women diagnosed with endometrial cancer between 1991 and 2000: (i) all 806 black women, (ii) 806 matched white women who appeared to have similar cancer at diagnosis, (iii) 806 other matched white women who appeared to have similar cancer at diagnosis and received similar surgical treatment, and (iv) all 13,756 white women. Both matched groups were matched for age, clinical stage or missing clinical stage, grade or missing grade, histology or missing histology, and the presence or absence of comorbid conditions; these variables are described in the top half of Table 1. In addition, the whites matched at treatment were matched for adequate surgery (AS = yes or no), surgeon type (GO = gynecological oncologist, GYN = gynecologist, Other = other surgeon), and an approximate measure of days from diagnosis to surgery, where surgeon type and days are defined only for the subset of women who received adequate surgery; these variables are described in the bottom half of Table 1.

In Table 1, black women were sicker at diagnosis than all white women. Black women tended to have more advanced stage of cancer, higher grade, and were more likely to present with a comorbid condition. In contrast, both matched groups of white women were similar to the black women in terms of the variables in the top half of Table 1. Black women received inferior medical care when compared to all white women and to the first matched group of white women who appeared similar in terms of the biological variables in Table 1. Black women were more likely to receive either no surgery or inadequate surgery (35% for blacks, 18% for similar whites matched at diagnosis, and 14% for the healthier group of

Table 1. Matching variables for all blacks, two nonoverlapping matched comparison groups of whites, and all whites. Numbers are percents, except as noted.

Matching variables at diagnosis	Black	Matched whites at diagnosis	Matched whites at treatment	All whites
<i>n</i>	806	806	806	13,756
Age (mean)	75	75	75	75
Stage 1	46	46	46	69
Stage 2	12	12	12	8
Stage 3	11	11	11	8
Stage 4	16	16	16	8
Stage Missing	14	14	14	7
Total %	100	100	100	100
Grade 1	14	14	14	32
Grade 2	21	21	22	33
Grade 3 or 4	38	38	37	24
Grade Missing	27	27	27	11
Total %	100	100	100	100
Adenocarcinoma	41	41	43	59
Endometroid	19	19	19	25
Papillary	23	23	22	10
Unknown	17	17	16	6
Total %	100	100	100	100
Diagnosis Year (mean)	1995	1995	1995	1995
Comorbidity %	43	42	43	26
Matching Variables at Treatment	Black	Matched Whites at Diagnosis	Matched Whites at Treatment	All Whites
Inadequate or No Surgery	35	18	35	14
Adequate Surgery (AS)	65	82	65	86
Total %	100	100	100	100
AS by GO Surgeon	28	28	28	24
AS by GYN Surgeon	25	43	26	50
AS by Other Surgeon	12	11	12	12
Inadequate or No Surgery	35	18	35	14
Total %	100	100	100	100
≤ 50 Days to AS	47	68	49	71
> 50 Days to AS	18	13	17	15
Inadequate or No Surgery	35	18	35	14
Total %	100	100	100	100

all whites.) In contrast, the second matched group of white women was similar to the black women in terms of both the biological variables in Table 1 and the treatment variables in Table 1.

Comparisons of survival in matched pairs are based on Albers (1988) rank test for randomly censored matched pairs, using his Wilcoxon-type scores. Largely because blacks were diagnosed with more advanced cancer, blacks typically survived for a much shorter time than the population of all whites (median 2.9 years for blacks, with 95% confidence interval [2.4, 3.3] versus 8.7 years for all whites, [8.4, 8.9]). White women matched for health at diagnosis survived somewhat longer than the black women (median 4.1 years, [3.5, 4.8], with Albers' one-sided p -value 0.037). Recall from Table 1 that this first matched group of white women was more likely to receive adequate surgery than the group of black women. When matched also for surgery, the second matched group of white women still had somewhat longer survival than the black women (median 3.6 years, [2.7, 4.3], with Albers' one-sided p -value 0.033). Survival in the two nonoverlapping matched groups of white women did not differ significantly, despite the higher frequency of adequate surgery in the first matched group (Albers' one-sided p -value 0.48).

In short, blacks are diagnosed with more advanced cancer, and this might account for much of the difference in survival. Blacks are less likely to receive adequate surgery than matched whites who were similar at diagnosis, but there is little indication that this difference in surgery is the cause of the remaining difference in survival, because whites who received similar surgical treatment still survived longer. These observations reflect the covariates in Table 1, and might be different with more detailed clinical data.

2.2 ADVANTAGES OF TAPERED MATCHING IN THIS EXAMPLE

The example in Section 2.1 illustrates several advantages of tapered matching.

Focus on the correct population. The black population is the natural one: (essentially) all blacks diagnosed with endometrial cancer in the SEER-Medicare cohort, and that population does not change as the analysis proceeds. All blacks are compared to whites who become more and more similar to the blacks, but the comparisons always refer to, and are weighted by, the entire actual population of blacks. The question, "How beneficial is adequate surgery?," is a different question if answered in the black population than in the entire population, which is mostly white and mostly diagnosed at an earlier, healthier stage. The analysis in Section 2.1 suggests substantial improvements in survival for blacks are mostly likely to be achieved through earlier diagnosis, despite the fact that there is also a disparity in the use of adequate surgery. See Dehejia and Wahba (1999) for discussion of the dangers of model-based adjustments, unaided by matching, which may not focus on the correct population.

Adjusting for interactions. In clinical data, the very meaning of one matching variable often changes depending upon the value of another matching variable, and for this reason it is important to control for interactions. The stage of endometrial cancer can be determined approximately by biopsy, more accurately as a by-product of adequate surgery, and still more accurately as a by-product of accurate surgery by a

gynecological oncologist whose specialty is gynecological cancer surgery. Because of this, the meaning of “stage 3” is different for a woman who had adequate surgery than for one who did not, and possibly different depending upon the surgeon type (Earle et al. 2006; Silber et al. 2007). Time to surgery and surgeon type in Table 1 have meaning only for a patient who has surgery. A “missing stage” is unlikely to be missing completely at random, as it would result if even a biopsy was not performed; see the appendix of Rosenbaum and Rubin (1984) for the relationship between matching and incomplete covariates. In Table 1, blacks lacking key clinical data are matched to whites lacking the same data, and whites matched at treatment are matched to blacks for “adequate surgery” and typically for surgeon type, so in these black-white pairs, important interactions are controlled.

The comparison is easy to understand. Although matching itself requires some technology, in the end, patients who are comparable in specific respects are compared in a straightforward way. Anyone can understand the degree and nature of the comparability displayed in Table 1, and the comparison of survival outcomes. If scientific findings are to affect policy, it is helpful to have an analysis that nonstatisticians can understand.

3. OPTIMAL TAPERED MATCHING USING THE ASSIGNMENT ALGORITHM

3.1 THE TAPERED MATCHING PROBLEM

The focal group consists of I individuals, $\alpha_1, \dots, \alpha_I$, and the comparison population consists of J individuals, $\gamma_1, \dots, \gamma_J$. In Section 2, $I = 806$ and $J = 13,756$. A total of $C \geq 2$ tapered, nonoverlapping comparison groups will be formed, $c = 1, \dots, C$, and α_i will be matched to $k_c \geq 1$ individuals in comparison group c such that each of the J individuals $\gamma_1, \dots, \gamma_J$ appears at most once among the $T = IK$ matched controls, where $K = \sum k_c$. In Section 2, $C = 2$, $k_1 = k_2 = 1$, $K = 2$, $T = IK = 1612$. Matching is possible only if $J \geq T$. Let $|H|$ denote the number of elements of a finite set H .

The variables used in forming comparison group c are different from the variables used in forming comparison group c' for $c' \neq c$, so the distance between α_i and γ_j changes with c . In Section 2, the variables used in forming comparison group $c = 1$ described the health of patients at diagnosis, whereas the variables used in forming group $c = 2$ described both the health of patients at diagnosis and their surgical treatment. Let $\lambda_{cij} \geq 0$ be the distance between α_i and γ_j when forming group c . In Section 2, λ_{1ij} is the distance between black woman α_i and white woman γ_j in terms of variables describing health at diagnosis, whereas λ_{2ij} describes the distance between these same two women in terms of health at diagnosis and surgical treatment. Two women, α_i and γ_j , might have been similar in terms of health at diagnosis, so λ_{1ij} is small, but they might have been treated very differently, so λ_{2ij} is large. An infinite distance, $\lambda_{cij} = \infty$, is used to prevent matching α_i and γ_j at level c of the taper. In Section 2, we insisted on an exact match on clinical stage at both levels of the taper, $c = 1, 2$, and an exact match on the binary indicator of adequate

surgery at level $c = 2$ of the taper, so $\lambda_{cij} = \infty$ if a_i and γ_j differed in these ways.

The problem is to determine the subset of $\{\gamma_1, \dots, \gamma_J\}$ to be matched to a_i at each level c of the taper. Let $\omega_{ci} \subset \{\gamma_1, \dots, \gamma_J\}$ be the set containing the k_c individuals matched to a_i at level c of the taper; so ω_{ci} identifies k_c controls, $|\omega_{ci}| = k_c$, and no control is used twice, $\omega_{ci} \cap \omega_{c'i'} = \emptyset$ if $i' \neq i$ or $c' \neq c$, and there are $T = IK$ distinct controls in all,

$$T = IK = \left| \bigcup_{i=1}^I \bigcup_{c=1}^C \omega_{ci} \right|.$$

Write Ω for the $I \times C$ array of sets ω_{ci} , $i = 1, \dots, I$, $c = 1, \dots, C$. The optimal tapered matching problem, Problem 1, is to construct Ω in such a way as to minimize the total of the distances λ_{cij} between focal group members, a_i , and their matched controls, γ_j , $j \in \omega_{ci}$.

Problem 1. An $I \times C$ array Ω of sets $\omega_{ci} \subset \{\gamma_1, \dots, \gamma_J\}$, $i = 1, \dots, I$, $c = 1, \dots, C$, such that $|\omega_{ci}| = k_c$, $\omega_{ci} \cap \omega_{c'i'} = \emptyset$ if $i' \neq i$ or $c' \neq c$ is called feasible. Find a feasible array Ω that minimizes $\lambda(\Omega) = \sum_{i=1}^I \sum_{c=1}^C \sum_{j \in \omega_{ci}} \lambda_{cij}$.

3.2 AN EQUIVALENT ASSIGNMENT PROBLEM

Define a matrix Δ with $T = IK$ rows and J columns as follows. The J columns of Δ correspond with $\gamma_1, \dots, \gamma_J$. The IK rows of Δ divide into I groups of K rows, one group for each a_i . The $K = \sum k_c$ rows in each of the I groups further divide into k_1 rows for the first level of the taper, $c = 1$, k_2 rows for the second level, $c = 2, \dots$, k_C rows for level C . Form Δ by placing λ_{cij} in the k_c rows

$$(i-1)K + \sum_{b=1}^{c-1} k_b + 1, \dots, (i-1)K + \sum_{b=1}^{c-1} k_b + k_c, \quad (3.1)$$

and column j , for $i = 1, \dots, I$, $c = 1, \dots, C$, $j = 1, \dots, J$, with $\sum_{b=1}^0 k_b = 0$.

Method 1: Solve the optimal assignment problem for the distance matrix Δ . For the resulting optimal assignment, define ω_{ci} to be the columns paired with the rows for a_i with the k_c indices (3.1), for $i = 1, \dots, I$, $c = 1, \dots, C$.

Proposition 2. The tapered matching method solves Problem 1.

Proof: Each $\beta(\cdot) \in \mathcal{B}$ determines a feasible Ω in Problem 1 using Method 1; then, the ω_{ci} so defined are disjoint because $\beta(\cdot)$ is an assignment and have $|\omega_{ci}| = k_c$ by construction. Conversely, if Ω is any feasible array of ω_{ci} in Problem 1, then Ω picks out $T = IK$ distinct columns of Δ from $\{\gamma_1, \dots, \gamma_J\}$ because the ω_{ci} do not overlap. For each c and i , there are $k_c!$ ways to assign the columns in ω_{ci} to the rows of Δ with indices (3.1), so each Ω corresponds with $\left(\prod_{c=1}^C k_c!\right)^I$ different assignments $\beta(\cdot) \in \mathcal{B}$ with the same cost $\delta\{\beta(\cdot)\}$ which equals $\lambda(\Omega)$. It follows that any minimum cost assignment $\beta(\cdot)$ corresponds with a minimum cost feasible array Ω . \square

In Hansen's (2007) `optmatch` package, when $k_1 = \dots = k_C > 1$, a smaller distance matrix with IC rows may be used in place of Δ in (3.1) with $IK \geq IC$ rows, by instructing

the package to match each row to $k_1 > 1$ controls. In R, an unnecessarily large distance matrix may lead to an inefficient use of memory, with slower performance.

3.3 DETAILS OF IMPLEMENTATION IN THE EXAMPLE

The optimization used Hansen's (2005, 2007) `optmatch` package in R. The two distances, λ_{cij} , $c = 1, 2$, were each formed using Mahalanobis metric matching within calipers on two propensity scores. The two distances, λ_{cij} , $c = 1, 2$, for matching at the time of diagnosis, $c = 1$, and at the time of treatment, $c = 2$, were defined as follows. For the diagnosis match, the probability that a patient is black rather than white was estimated from the variables in the top half of Table 1 using a logit model; for matching purposes, this is similar to a propensity score. For the treatment match, the probability that a patient is black rather than white was estimated from all of the variables in Tables 1 using a logit model. For discussion of the consequences of using an estimated propensity score, see Hirano et al. (2003). In both matches, we insisted on an exact match on clinical stage or stage missing. This can be implemented with infinite distances, but in fact it is more efficient in time and space to divide the matching problem into five separate problems, one for each level of stage. The distance λ_{1ij} was calculated as the Mahalanobis distances (Rubin 1979) for the variables in the top of Table 1, and λ_{2ij} was calculated as the Mahalanobis distances for all the variables in Tables 1, with certain penalties added. Specifically, a penalty of 10^6 was added to λ_{2ij} if α_i and γ_j differed on the binary indicator of adequate surgery; this ensured that in the treatment match, $c = 2$, patients were exactly matched for adequate surgery. In addition, penalties of 10^3 were added to λ_{cij} if α_i and γ_j differed on propensity score c by more than one fifth of a standard deviation of the propensity score c , $c = 1, 2$, and a penalty of 10^3 was added if α_i and γ_j differed on the binary indicator for comorbidity.

This use of penalties is standard in optimization: it replaces a constraint by a change in the objective function. With a sufficiently large penalty, if the constraint can be respected, it will be, and if the constraint cannot be respected in every case, it will be respected as often as possible. In Table 1, in the match at diagnosis, the comorbidity constraint was respected in $796/806 = 98.8\%$ of pairs. In the match at treatment, there are $805/806 = 99.9\%$ pairs matched for comorbidity.

3.4 CHOICE OF DISTANCE

In tapered matching, two or more definitions of "distance," say λ_{1ij} and λ_{2ij} , are added together in $\lambda(\Omega)$. These different distances need to be such that they can reasonably be added together. Use of the Mahalanobis distance, with or without added penalties, is attractive in this context. This brief section explains why.

Suppose that the covariates for the $I + J$ individuals $\{\alpha_1, \dots, \alpha_I, \gamma_1, \dots, \gamma_J\}$ are in a matrix $\mathbf{X} = (\mathbf{X}_1, \mathbf{X}_2, \dots, \mathbf{X}_C)$ with $I + J$ rows and $L = \ell_1 + \dots + \ell_C$ columns, where level $c = 1$ of the taper will match on the ℓ_1 variables in $\mathbf{V}_1 = \mathbf{X}_1$, level $c = 2$ of the taper will match on the $\ell_1 + \ell_2$ variables in $\mathbf{V}_2 = (\mathbf{X}_1, \mathbf{X}_2)$, and generally, level c of the taper will match on the $\ell_1 + \dots + \ell_c$ variables in $\mathbf{V}_c = (\mathbf{X}_1, \mathbf{X}_2, \dots, \mathbf{X}_c)$, for $c = 1, \dots, C$. Here, row i of \mathbf{V}_c , say \mathbf{v}_{ci} , describes α_i , $i = 1, \dots, I$, while row $I + j$ of \mathbf{V}_c , say $\mathbf{v}_{c,I+j}$, describes γ_j , $j = 1, \dots, J$ and each is of dimension $\ell_1 + \dots + \ell_c$.

Suppose that λ_{cij} is the sample Mahalanobis distance between α_i and γ_j computed from $\mathbf{V}_c = (\mathbf{X}_1, \mathbf{X}_2, \dots, \mathbf{X}_c)$, that is, $\lambda_{cij} = (\mathbf{v}_{ci} - \mathbf{v}_{c,I+j}) \widehat{\Sigma}_c^{-1} (\mathbf{v}_{ci} - \mathbf{v}_{c,I+j})^T$, where $\widehat{\Sigma}_c$ is the sample covariance matrix computed from \mathbf{V}_c . The Mahalanobis distances are invariant under affine transformations of the data matrix. In particular, if Gram–Schmidt orthogonalization (Rao 1973, sec. 1b.2ix) were applied to the columns of $\mathbf{X} = (\mathbf{X}_1, \mathbf{X}_2, \dots, \mathbf{X}_C)$, and the resulting columns were standardized to mean zero and unit variance, and the Mahalanobis distances, the λ_{cij} 's, were recomputed from these uncorrelated, standardized variables, then all $C \times I \times J$ distances λ_{cij} would be unchanged. In other words, without loss of generality, one may view all $C \times I \times J$ Mahalanobis distances λ_{cij} as having been computed from uncorrelated variables with sample mean zero and sample standard deviation one. For such variables, the Mahalanobis distance λ_{cij} is the sum of $\ell_1 + \dots + \ell_c$ squared differences in the variables describing α_i and γ_j . In other words, in the criterion $\lambda(\Omega)$, a difference of one standard deviation in one of these standardized variables contributes a value of one no matter where it occurs. The first variable in Table 1 is age. If the standard deviation of age in this Medicare population were five years, and if α_i were 70 and γ_j were 75 then α_i and γ_j would differ by one standard deviation. Speaking informally, the two Mahalanobis distances count this five year difference in age between α_i and γ_j as equally serious for the match at diagnosis and the match at treatment, even though there are more variables in the match at treatment, so the distances tend to be larger in total at treatment. Recall that Gram–Schmidt orthogonalization makes the columns of \mathbf{X} orthogonal by, essentially, finding the residuals when column k of \mathbf{X} is regressed on columns 1, $\dots, k-1$, for $k = 2, \dots, L$. Speaking somewhat more formally, the Mahalanobis distance between α_i and γ_j at treatment, namely λ_{2ij} , equals the Mahalanobis distance between α_i and γ_j at diagnosis, namely λ_{1ij} , plus the Mahalanobis distance between α_i and γ_j on the residuals of the treatment variables after regression on the diagnostic variables. In consequence, the difference between α_i and γ_j on the diagnostic variables in Table 1 counts the same, namely λ_{1ij} , at both levels of the taper, but there is an additional contribution from the treatment variables, so $\lambda_{2ij} \geq \lambda_{1ij}$. In words, the Mahalanobis distance does not regard mismatches of a certain magnitude on the diagnostic variables to be more or less important when matching also on the surgical variables, and this is appropriate if both groups are to be well-matched in Table 1. Similar considerations apply to the large penalties, whether they are added because of violations of the propensity score caliper or because of mismatches on a specific variable such as comorbidity.

ACKNOWLEDGMENTS

This work was supported by grants from the Measurement, Methodology and Statistics Program of the U.S. National Science Foundation and from the U.S. National Cancer Institute.

[Received July 2007. Revised April 2008.]

REFERENCES

Albers, W. (1988), "Combined Rank Tests for Randomly Censored Paired Data," *Journal of the American Statistical Association*, 83, 1159–1161.

- Augurzyk, B., and Kluge, J. (2007), "Assessing the Performance of Matching Algorithms When Selection into Treatment is Strong," *Journal of Applied Econometrics*, 22, 533–557.
- Bergstralh, E. J., Kosanke, J. L., and Jacobsen, S. L. (1996), "Software for Optimal Matching in Observational Studies," *Epidemiology*, 7, 331–332. SAS code is available at online <http://www.mayo.edu/hsr/sasmac.html>.
- Bertsekas, D. P. (1981), "A New Algorithm for the Assignment Problem," *Mathematical Programming*, 21, 152–171. Fortran code is available online at <http://web.mit.edu/dimitrib/www/home.html>.
- Carpaneto, G., and Toth, P. (1980), "Algorithm 548: Solution of the Assignment Problem [H]," *ACM Transactions on Mathematical Software*, 6, 104–111. Fortran code is available at <http://portal.acm.org>.
- Dehejia, R., and Wahba, S. (1999), "Causal Effects in Nonexperimental Studies: Reevaluating the Evaluation of Training Programs," *Journal of the American Statistical Association*, 94, 1053–1062.
- Dell'Amico, M., and Toth, P. (2000), "Algorithms and Codes for Dense Assignment Problems: The State of the Art," *Discrete Applied Mathematics*, 100, 17–48.
- Earle, C. C., Schrag, D., Neville, B. A., et al. (2006), "Effect of Surgeon Specialty on Processes of Care and Outcomes for Ovarian Cancer Patients," *Journal of the National Cancer Institute*, 98, 172–180.
- Hansen, B. B. (2004), "Full Matching in an Observational Study of Coaching for the SAT," *Journal of the American Statistical Association*, 99, 609–618.
- Hansen, B. B., and Klopfer, S. O. (2006), "Optimal Full Matching and Related Designs via Network Flows," *Journal of Computational and Graphical Statistics*, 15, 609–627.
- Hansen, B. B. (2007), "Optmatch: Flexible, Optimal Matching for Observational Studies," *R News*, 7, 18–24.
- Hirano, K., Imbens, G. W., and Ridder, G. (2003), "Efficient Estimation of Average Treatment Effects using the Estimated Propensity Score," *Econometrica*, 71, 1161–1189.
- Lu, B. (2005), "Propensity Score Matching with Time-Dependent Covariates," *Biometrics*, 61, 721–728.
- Papadimitriou, C., and Steiglitz, K. (1982), *Combinatorial Optimization*, Englewood Cliffs, NJ: Prentice-Hall.
- R Development Core Team (2007), *R: A Language and Environment for Statistical Computing*, Vienna: R Foundation, <http://www.R-project.org>.
- Rao, C. R. (1973), *Linear Statistical Inference and Applications*, New York: Wiley.
- Rosenbaum, P. R. (1989), "Optimal Matching in Observational Studies," *Journal of the American Statistical Association*, 84, 1024–1032.
- (1991), "A Characterization of Optimal Designs for Observational Studies," *Journal of the Royal Statistical Society, Series B*, 53, 597–610.
- Rosenbaum, P., and Rubin, D. B. (1984), "Reducing Bias in Observational Studies using Subclassification on the Propensity Score," *Journal of the American Statistical Association*, 79, 516–524.
- 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, 75–83.
- Rubin, D. B. (1979), "Using Multivariate Matched Sampling and Regression Adjustment to Control Bias in Observational Studies," *Journal of the American Statistical Association*, 74, 318–328.
- Silber, J. H., Rosenbaum, P. R., Trudeau, M. E., Even-Shoshan, O., Chen, W., Zhang, X., and Mosher, R. E. (2001), "Multivariate Matching and Bias Reduction in the Surgical Outcomes Study," *Medical Care*, 39, 1048–1064.
- Silber, J. H., Rosenbaum, P. R., Trudeau, M. E., Chen, W., Zhang, X., Lorch, S., Rapaport-Kelz, R., Mosher, R. E., Even-Shoshan, O. (2005), "Preoperative Antibiotics and Mortality in the Elderly," *Annals of Surgery*, 242, 107–114.
- Silber, J. H., Rosenbaum, P. R., Polsky, D. E., Ross, R. N., Even-Shoshan, O., Schwartz, J. S., Armstrong, K. A., and Randall, T. C. (2007), "Do Ovarian Cancer Outcomes Differ by Who Gives the Chemotherapy?" *Journal of Clinical Oncology*, 25, 1169–1175.
- Smith, H. (1997), "Matching with Multiple Controls to Estimate Treatment Effects in Observational Studies," *Sociological Methodology* 27, 325–353.