
Multiple Linear Regression Viewpoints

A Publication sponsored by the American Educational
Research Association's Special Interest Group on
Multiple Linear Regression: The General Linear Model

MLRV

Volume 35 • Number 1 • Fall 2009

Table of Contents

An Application of a New Multivariate Resampling Method to Multiple Regression	1
Haiyan Bai	University of Central Florida
Wei Pan	University of Cincinnati
Interaction Effects: Centering, Variance Inflation Factor, and Interpretation Issues	6
Cecil Robinson	University of Alabama
Randall E. Schumacker	University of Alabama
Using Linear Regression and Propensity Score Matching to Estimate the Effect of Coaching on the SAT	12
Ben Domingue	University of Colorado
Derek C. Briggs	University of Colorado
Hierarchically Ordered Instructional Activities for the Development of Clinically Relevant Competencies: A Sequential Path Model	30
Frank J. Papa	University of North Texas Health Science Center
Randall E. Schumacker	University of Alabama

Multiple Linear Regression Viewpoints

Editorial Board

Randall E. Schumacker, Editor
University of Alabama

T. Mark Beasley, Associate Editor
University of Alabama at Birmingham

Isadore Newman, Editor Emeritus
Florida International University

Gordon P. Brooks (2008-2010) Ohio University
Daniel J. Mundfrom (2008-2010) Northern Colorado University
Bruce G. Rogers (2008-2010) University of Northern Iowa
Thomas Smith (2008-2010) Northern Illinois University
Susan Tracz (2008-2010) California State University, Fresno
David Walker (2008-2010) Northern Illinois University

Multiple Linear Regression Viewpoints (ISSN 0195-7171) is published by the AERA Special Interest Group on Multiple Linear Regression: General Linear Model through the **University of Alabama at Birmingham**.

Subscription and SIG membership information can be obtained from:
Cynthia Campbell, MLR:GLM/SIG Executive Secretary
Department of Educational Technology, Research & Assessment
Northern Illinois University
DeKalb, IL 60115-2854.
ccampbell@niu.edu

MLRV abstracts appear in CIJE, the ERIC system, and microform copies are available from University Microfilms International, 300 North Zeeb Road, Ann Arbor, MI 48106. *MLRV* is listed in the *EBSCO Librarians Handbook*.

An Application of a New Multivariate Resampling Method to Multiple Regression

Haiyan Bai

University of Central Florida

Wei Pan

University of Cincinnati

This study presents a new multivariate resampling method to improve the performance of multiple regression with small samples. The kernel resampling technique (KRT) is utilized in the multivariate resampling procedure to draw random resamples with random noises, which facilitates obtaining more accurate parameter estimates and their standard errors in multiple regression. The findings from an empirical example suggest that the statistical performance of multiple regression is improved through the KRT technique.

Multiple regression is one of the popular statistical methods; however, when it is applied to small samples, it encounters problems related to the accuracy of statistical estimation (Allison, 1999). Resampling method has been implemented to solve the small sample problems (Bai & Pan, 2008; Davison & Hinkley, 1997; Efron & Tibshirani, 1998). The bootstrap as the most popular resampling method has been applied in regression analysis since Efron's pioneer work in 1979. The bootstrap is an effective tool to solve the small sample problems, and comprehensive applications of the bootstrap to regression models have also been developed by researchers (e.g., Bickel & Freedman, 1981, 1983; Freedman, 1981; Peters & Freedman, 1984; Shao, 1988; Weber, 1984; Wu, 1986); however, the bootstrap standard errors tend to be biased downward in regression analysis when applying to a small sample (Peters & Freedman, 1984). Therefore, managing small sample problems in regression still remains a pertinent issue. A study on improving the statistical performance of multiple regression with small samples would significantly contribute to the literature in both methodological research of small sample issues and applied research using multiple regression with small samples.

The purpose of this present study is to introduce a new multivariate resampling method, the kernel resampling technique (KRT), to tackle the problems in multiple regression with small samples. Specifically, the present study (a) introduces the procedure of KRT, (b) examines the performance of KRT in multiple regression through an empirical example, and (c) compares the performance of KRT in multiple regression with that of the bootstrap in terms of estimation bias and standard errors.

Kernel Resampling Technique

Kernel Resampling Technique (KRT) is a new resampling method which uses kernel smoothing technique to capture the shape of the empirical sample distributions and sampling from the neighborhoods of the original sample. The kernel technique uses kernel probability density estimation to map the original linear or non-linear observations into a higher-dimensional space, where the linear classifier is subsequently used to solve problems (Aizerman, Braverman, & Rozonoer, 1964). The kernel probability density estimation has gained popularity, especially for dealing with nonparametric issues (Towers, 2002).

The kernel technique has been used in the bootstrap for smoothing the bootstrap distribution (Efron & Tibshirini, 1998; Silverman & Young, 1987). However, it is worth noting that there are two key points when using the kernel technique in the smooth bootstrap: (1) kernels are used *after* the bootstrap resampling to smooth the bootstrap distribution, and (2) the bandwidths of the kernels are not specifically defined for different data in the smooth bootstrap. With regards to the second point, finding an optimal bandwidth for the bootstrap smoothing procedure is a statistically and technically challenging task for many researchers and statisticians (Silverman & Young, 1987).

Regarding the issues of the basic and smooth bootstrap, this present study presents a new multivariate kernel resampling technique for improving the performance of multiple regression with small samples because kernel technique has been proved remarkably successful for standard classification and regression problems (Schölkopf & Smola, 2002). KRT is a new resampling method where the overall procedure seems similar to that of the bootstrap, but KRT, by design, radically differs from the bootstrap in three-fold: (a) The bootstrap samples are randomly drawn from the exact original small sample data with replacement in the basic bootstrap, whereas the KRT samples are each randomly drawn from a *neighborhood* of a data point in the original small sample; (b) The kernel technique is used to select resamples with random noises instead of being used to smooth the resampling distributions as does the

smooth bootstrap; and (c) KRT has a fixed bandwidth for the kernel used in the resampling procedures whereas the smooth bootstrap does not.

KRT utilizes Gaussian kernels (Silverman, 1986; Simonoff, 1996), the most commonly-used kernel technique (Yip, Ahmad, & Pong, 1999), to capture the underlying distribution of the given multivariate small sample data. The Gaussian kernel bandwidth is determined to produce an asymptotically optimal bandwidth minimizing the *mean integrated square error* (MISE; Silverman, 1986).

Specifically, let $\mathbf{X}_1, \dots, \mathbf{X}_n$ be the given multivariate small sample data or n vectors from a d -dimensional space R^d , where d is the number of variables and n is the sample size. The KRT procedure is as follows:

Step 1. Define n multivariate Gaussian kernels as $K_i(\mathbf{x}) \sim N_d(\mathbf{X}_i, \mathbf{H}_0^2)$, $i = 1, \dots, n$, where the mean vector \mathbf{X}_i is the i th multivariate observation in the multivariate small sample and the random noise \mathbf{H}_0 can be determined by the optimal bandwidth MISE (Silverman, 1986, p. 87; Simonoff, 1996, p. 105):

$$\mathbf{H}_0 = \left(\frac{4}{d+2} \right)^{1/(d+4)} \Sigma^{1/2} n^{-1/(d+4)}, \quad (1)$$

where Σ is the population covariance matrix and it can be estimated by \mathbf{S} , a sample covariance matrix of $\mathbf{X}_1, \dots, \mathbf{X}_n$.

Step 2. Draw n multivariate random observations, $\mathbf{X}^*_1, \dots, \mathbf{X}^*_n$, each from one multivariate Gaussian kernel $K_i(\mathbf{x})$ ($i = 1, \dots, n$). The n multivariate random observations are defined as a multivariate KRT sample. According to Silverman (1986, eq. 4.7, p. 78) and Simonoff (1996, eq. 4.5, p. 102), the multivariate KRT sample has an estimated multivariate density function as follows:

$$\hat{f}(\mathbf{x}) = \frac{1}{n|\mathbf{H}|} \sum_{i=1}^n k_d[\mathbf{H}^{-1}(\mathbf{x} - \mathbf{X}_i)], \quad (2)$$

where $k_d(\mathbf{u}) \sim N_d(\mathbf{0}, \mathbf{I})$ and has the following distributional relationship with the multivariate Gaussian kernel $K_i(\mathbf{x})$, $i = 1, \dots, n$:

$$\begin{aligned} k_d(\mathbf{u}) &= (2\pi)^{-d/2} \exp\left(-\frac{1}{2}\mathbf{u}^T\mathbf{u}\right) \\ &= (2\pi)^{-d/2} \exp\left[-\frac{1}{2}(\mathbf{x} - \mathbf{X}_i)^T\mathbf{H}^{-2}(\mathbf{x} - \mathbf{X}_i)\right] \\ &= K_i(\mathbf{x}). \end{aligned} \quad (3)$$

Step 3. Conduct multiple regression over the multivariate KRT sample, $\mathbf{X}^*_1, \dots, \mathbf{X}^*_n$, to obtain an estimate of a parameter of interest.

Step 4. Repeat Steps 2 to 3 k times, where k is called the KRT resampling parameter, to obtain k parameter estimates that comprise a sampling distribution of the parameter of interest.

Step 5. Evaluate the performance of KRT in regression analysis based on the sampling distribution such as estimation bias and standard errors.

The above procedure has been written into a SAS macro program and is available through the author. The “plug-in principle” (Efron & Tibshirani, 1998, p.35) makes the application of KRT very simple and straightforward because the above procedure does not require researchers to modify the bandwidth of the kernel to obtain KRT samples. Therefore, KRT is methodologically comparable to the bootstrap but the application of KRT is simpler than that of the bootstrap.

In the next section, an empirical example is presented for illustrating the application of the KRT procedure to multiple regression and evaluating the statistical performance of KRT in multiple regression with a small sample.

An Empirical Study of KRT in Multiple Regression Analysis

The Cement Hardening Data and Regression Model

The famous small sample of the Cement Hardening Data (CH) (Hjorth, 1994, p. 31) was used to study the performance of the application of the KRT procedure in multiple regression while comparing

Table 1. *The Cement Hardening Data*

x_1	x_2	x_3	x_4	y
7	26	6	60	78.5
1	29	15	52	74.3
11	56	8	20	104.3
11	31	8	47	87.6
7	52	6	33	95.9
11	55	9	22	109.2
3	71	17	6	102.7
1	31	22	44	72.5
2	54	18	22	93.1
21	47	4	26	115.9
1	40	23	34	83.8
11	66	9	12	113.3
10	68	8	12	109.4

Note. x_1 = amount of tricalcium aluminate, $3\text{CaO} \cdot \text{Al}_2\text{O}_3$; x_2 = amount of tricalcium silicate, $3\text{CaO} \cdot \text{SiO}_2$; x_3 = amount of calcium aluminum ferrate, $4\text{CaO} \cdot \text{Al}_2\text{O}_3 \cdot \text{Fe}_2\text{O}_3$; x_4 = amount of dicalcium silicate, $2\text{CaO} \cdot \text{SiO}_2$; y (response) = heat evolved in calories per gram of cement.

The CH data (Table 1) with 13 observations depict the hardening of the cement and the heat evolved during the first 180 days after addition of water. x_1 , x_2 , x_3 , and x_4 were linearly dependent predictors of the amounts of different components of chemicals, and the response variable y was the heat produced. Because Hjorth's (1994) linear regression model (4) was proven to have a good fit with inclusion of all the four predictors presented, it was used for both the KRT samples and the bootstrapping observations.

$$y_i = \beta_0 + \beta_1 x_{i1} + \beta_2 x_{i2} + \beta_3 x_{i3} + \beta_4 x_{i4} + \varepsilon_i, i = 1, 2, \dots, n. \quad (4)$$

Comparisons of Estimation Accuracy for the Empirical data

Table 2 shows the ordinary least-squares (OLS) estimates ($\hat{\beta}_0$, $\hat{\beta}_1$, $\hat{\beta}_2$, $\hat{\beta}_3$, $\hat{\beta}_4$) from the original small sample, the estimates from 200 KRT samples and 200 bootstrap samples and their estimated standard errors $\overline{se}(\hat{\beta}_i)$, $i = 1, \dots, 4$, from the regression model (4).

From Table 2 we can see that the estimates from KRT were systematically closer to the estimates from OLS than the estimates from the bootstrap; and the standard errors estimated from KRT were also systematically smaller than those from both the bootstrap and OLS. In addition, it is clear that the approximate biases of the KRT estimates were systematically smaller than the biases of the bootstrap estimates.

Table 3 shows the model fit indices from all the methods. We can see from Table 3 that the root mean squared error (RMSE) estimate from the KRT procedure was also close to the estimate from the OLS method, but the RMSE estimate from the bootstrap was still downward biased for the analysis on the empirical data. The multiple correlation R^2 from the KRT procedure is also comparable to that from both the bootstrap and OLS method (see Table 3).

Results and Discussions

The present study presents the new multivariate resampling method, KRT, for obtaining more accurate estimates with reasonable standard errors in multiple regression analysis with small samples. Unlike the smooth bootstrap, which draws resamples from the smoothed distribution of the bootstrap data, KRT draws the resamples from the neighborhoods of the original data with a fixed but optimal bandwidth. As such, of KRT has the following advantages: (a) the resample distribution strictly follows the original sample distribution, (b) the sampling distribution is not artificially modified, and (c) there is no need for researchers to consider the kernel bandwidth.

The findings from the applications of KRT in multiple regression to the empirical data suggest that the KRT procedure outperformed other methods in terms of the accuracy of the estimation of regression coefficients, estimation bias, and standard errors, comparing with the OLS method on the original small

Table 2. Estimates of Regression Analysis on the Cement Hardening Data

Parameter	OLS Small N		Bootstrap ($B = 200$)				KRT ($k = 200$)			
	Est.	SE	Est.	Approx Bias	SE	Bias-Corrected Est.	Est.	Approx Bias	SE	Bias-Corrected Est.
β_0	62.41	70.07	74.00	11.59	117.10	50.81	61.10	-1.31	28.10	63.71
β_1	1.55	0.74	1.46	-0.09	1.16	1.64	1.57	0.02	0.30	1.53
β_2	0.51	0.72	0.39	-0.12	1.22	0.63	0.52	0.01	0.29	0.50
β_3	0.10	0.75	0.01	-0.09	1.20	0.20	0.12	0.02	0.30	0.08
β_4	-0.14	0.71	-0.27	-0.13	1.21	-0.02	-0.13	0.01	0.28	-0.16

Note. Est. = Estimate. Approx = Approximate.

Table 3. Model Fit Summary for the Cement Hardening Data

Parameter	OLS	Bootstrap ($B = 200$)			KRT ($k = 200$)			
	Small N	Est.	Approx Bias	SE	Est.	Approx Bias	SE	
RMSE	2.45	1.84	-0.60	0.55	2.59	0.15	0.28	
R^2	0.98	0.99	0.01	0.01	0.98	0.00	0.01	

Note. Est. = Estimate. Approx = Approximate.

sample and the bootstrap results. The results from this study also support the use of KRT as a viable alternative to improving the performance of multiple regression analysis with small samples. This current study suggests that KRT can be a useful tool for researchers to conduct multiple regression analysis when only small samples are available. It will help researchers draw more valid statistical inference than using the original small sample.

The advantage of KRT concerns the resampling procedure's simplicity and efficiency. Compared to the bootstrap, KRT obtains comparable or more accurate estimates, but does not require researchers to modify the complicated resampling procedures. The results from this study indicated that the KRT procedure has overcome the two major limitations that the bootstrap method encountered. First, KRT obtains independent resamples through sampling from the neighborhoods of the data points, which solves the lack of independent observations of the basic bootstrap resamples. Secondly, the KRT procedure practically advances the smooth bootstrap by using fixed optimal bandwidth for the kernel procedure instead of requiring researchers to customize the optimal bandwidth to their data. The simplicity of KRT will help improve the practical applications of resampling method in the real research and promote the use of resampling method in the computer age.

In the present study, we only explored the statistical performance of the application of the KRT procedure to the multiple regression in terms of the estimation of regression coefficients and model fit indices. Even though we have used the population parameters to verify the smaller biases for the estimations from the KRT procedure for the both estimates and standard errors, significance tests and the confidence intervals are desirable for further research.

References

- Aizerman, M., Braverman, E., & Rozonoer, L. (1964). Theoretical foundations of the potential function method in pattern recognition learning. *Automation and Remote Control*, 25, 821-837.
- Allison, P. D. (1999). *Multiple regression: A primer*. Thousand Oaks, CA: Pine Forge Press.
- Bai, H., & Pan, W. (2008). Resampling methods revisited: Advancing the understanding and applications in educational research. *International Journal of Research & Method in Education*, 31(1), 45-62.
- Bickel, P. J., & Freedman, D. A. (1981). Some asymptotic theory for the bootstrap. *The Annals of Statistics*, 9, 1196-1217.
- Bickel, P. J., & Freedman, D. A. (1983). Bootstrapping regression models with many parameters. In P. J. Bickel, K. Doksum, & J. L. Hodges (Eds.), *Festschrift for Erich Lehmann* (pp. 28-48). Belmont, CA: Wadsworth.
- Chernick, M. R. (1999). *Bootstrap methods: A practitioner's guide*. New York: Wiley and Sons, Inc.

- Davison, A. C., & Hinkley, D. V. (1997). *Bootstrap methods and their application*. New York: Cambridge University Press.
- Efron, B. (1979). Bootstrap methods: Another look at the jackknife. *The Annals of Statistics*, 7(1), 1-26.
- Efron, B. (1982). The jackknife, the bootstrap, and other resampling plans. *CBMS-NSF Regional Conference Series in Applied Mathematics* (No. 38). Philadelphia: SIAM.
- Efron, B., & Tibshirani, R. J. (1998). *An introduction to the bootstrap*. New York: Chapman & Hall/CRC.
- Freedman, D. A. (1981). Bootstrapping regression models. *The Annals of Statistics*, 9(6), 1218-1228.
- Hjorth, J. S. U. (1994). *Computer intensive statistical methods*. London: Chapman & Hall.
- Peters, S. C., & Freedman, D. A. (1984). Some notes on the bootstrap in regression problems. *Journal of Business and Economic Statistics*, 2, 401-409.
- SAS Institute Inc. (2008). *Sample 24982: Jackknife and bootstrap analyses*. Retrieved May 2, 2008, from <http://support.sas.com/kb/24/982.html>
- Schölkopf, B., & Smola, A. J. (2002). *A short introduction to learning with kernels*. New York: Springer-Verlag.
- Shao, J. (1988). On resampling methods for variance and bias estimation in linear models. *The Annals of Statistics*, 16, 986-1008.
- Shimabukuro, F. I., Lazar, S., Dyson, H. B., & Chernick, M. R. (1984). A quasi-optical method for measuring the complex permittivity of materials. *IEEE Transactions on Microwave Theory and Techniques*, 32, 659-665.
- Silverman, B. W. (1986). *Density estimation for statistics and data analysis*. New York: Chapman and Hall.
- Silverman, B. W., & Young, G. A. (1987). The bootstrap: To smooth or not to smooth? *Biometrika*, 74, 469-79.
- Simonoff, J. S. (1996). *Smoothing methods in statistics*. New York: Springer.
- Towers, S. (2002). Kernel probability density estimation methods. *Proceedings on Advanced Statistical Techniques in Particle Physics*, Durham, England.
- Weber, N. C. (1984). On resampling techniques for regression models. *Statistics & Probability Letters*, 2, 275-278.
- Wu, C. F. J. (1986). Jackknife bootstrap and other resampling plans in regression analysis (with discussion). *The Annals of Statistics*, 14, 1261-1350.
- Yip, H. M., Ahmad, I., & Pong, T. C. (1999). An efficient parallel algorithm for computing the Gaussian convolution of multi-dimensional image data. *Journal of Supercomputing*, 14, 233-255..

Send correspondence to: Haiyan Bai
 University of Central Florida
 Email: hbai@mail.ucf.edu

Interaction Effects: Centering, Variance Inflation Factor, and Interpretation Issues

Cecil Robinson

Randall E. Schumacker

University of Alabama

Research hypotheses that include interaction effects should be of more interest to educational researchers, especially since issues related to centering and interpretation of the variance inflation factor have been introduced. The purpose of this paper was to examine interaction effects in the context of centered versus uncentered variables and the variance inflation factor, especially upon the interpretation of interaction effects. Results indicated that centering of variables was required when examining interaction effects, uncentered variables impacted the variance inflation factor values, and separate regression equations have important interpretation outcomes in the presence of non-significant interaction effects.

Historically, hypotheses that specify testing interaction effects before examining main effects have appeared under the framework of analysis of variance. In the 1960's with the emergence of multiple regression, coding for interaction effects was introduced. Faculty who taught multiple regression therefore usually included instruction on dummy coding to obtain a test of interaction effects (Fox, 1997). Today, depending upon the textbook used, analysis of variance with A x B interaction effect may be covered without any corresponding interaction effect presentation given for multiple regression (Hinkle, Wiersma, & Jurs, 1998).

Much of the published research literature seems to only examine main effects or linear effects. In practice, A x B interactions are only found in a few published journal articles, A x B x C interactions, are less common, and A x B x C x D interactions are even more scarce. Only a few 5-way interactions have ever been published. The reason is that such higher level interaction effects are extremely difficult to interpret. Interaction effects that are categorical in nature, involve multiplicative continuous variables, or hypothesize quadratic or cubic terms are rare (Schumacker & Marcoulides, 1998).

Research hypotheses that include interaction effects should be of more interest to educational researchers, especially since issues related to centering (Aiken & West, 1991) and interpretation of the variance inflation factor (Freund, Littell, and Creighton, 2003) have been introduced. The purpose of this paper is to examine interaction effects in the context of centered versus uncentered variables and the variance inflation factor, especially upon the interpretation of interaction results.

Theoretical Framework

The effects of predictor scaling on coefficients of regression equations (centered versus uncentered solutions and higher order interaction effects (3-way interactions; categorical by continuous effects) has thoughtfully been covered by Aiken and West (1991). Their example illustrates that considerable multicollinearity is introduced into a regression equation with an interaction term when the variables are not centered. The variance inflation factor should detect the degree of multicollinearity when variables are uncentered (Freund, Little, & Creighton, 2003). The variance inflation factor as a measure of the degree of multicollinearity however has not been examined in context with centered versus uncentered variables in a regression equation containing interaction effects.

Centering

Centering is defined as subtracting the mean (a constant) from each score, X, yielding a centered score. Aiken & West (1991) demonstrated that using other transformations, additive constant, or uncentered scores can have a profound effect on interaction results. Regression with higher order terms has covariance between interaction terms (XZ) and each component (X and Z) depends in part upon the means of the individual predictors. Rescaling, changes the means, thus changes the predictor covariance, yielding different regression weights for the predictors in the higher order function. Centering is therefore an important step when testing interaction effects in multiple regression to obtain a meaningful interpretation of results.

Centering the variables places the intercept at the means of all the variables. A regression equation with an intercept is often misunderstood in the context of multicollinearity. The intercept is an estimate of the response at the origin where all independent variables are zero, thus inclusion of the intercept in the

study of collinearity is not of much interest. When variables have been centered, the intercept has no effect on the collinearity of the other variables (Belsley, Kuh, and Welsch, 1980).

Centering is also consistent with the computation of the variance inflation factor (VIF) and therefore it is suggested that VIF be computed only after first centering variables (Freund, Littell, and Creighton, 2003). Centered variables have low intercorrelation, while uncentered variables have higher intercorrelation, thus higher collinearity. The variance inflation factor is therefore an important part of examining interaction effects in multiple regression.

Variance Inflation Factor

When a full regression model is specified, multicollinearity amongst the predictor variables is possible. Multicollinearity can inflate the variance amongst the variables in the model. These inflated variances are problematic in regression because some variables add very little or even no new and independent information to the model (Belsley, Kuh & Welsch, 1980). Although Schroeder, Sjoquist and Stephen (1986) assert that there is no statistical test that can determine whether or not multicollinearity is a problem, there are ways for detecting multicollinearity (Berry & Feldman, 1985). For example, the variance inflation factor (VIF) can detect the degree of multicollinearity when variables are uncentered (Freund, Littell & Creighton, 2003). Stine (1985) also suggested a graphical approach to detecting VIF.

VIF measures the impact of multicollinearity among the X's in a regression model on the precision of estimation. It expresses the degree to which multicollinearity amongst the predictors degrades the precision of an estimate. VIF is a statistic used to measure possible multicollinearity amongst the predictor or explanatory variables. VIF is computed as $(1/(1-R^2))$ for each of the $k - 1$ independent variable equations. For example, given 4 independent predictor variables, the independent regression equations are formed by using each $k-1$ independent variable as the dependent variable:

$$\begin{aligned} X_1 &= X_2 X_3 X_4 \\ X_2 &= X_1 X_3 X_4 \\ X_3 &= X_1 X_2 X_4 \end{aligned}$$

Each independent variable model will return an R^2 value and VIF value. The term to exclude in the model is then based on the value of VIF. If X_j is highly correlated with the remaining predictors, its variance inflation factor will be very large. A general rule is that the VIF should not exceed 10 (Belsley, Kuh, & Welsch, 1980). When X_j is orthogonal to the remaining predictors, its variance inflation factor will be 1.

Method and Procedures

The rationale for the data analysis was that three concepts: self-efficacy (Bandura 1997), hope (Snyder, 1995) and optimism (Scheier & Carver, 1992); comprise a cognitive set that form a belief system which influenced academic achievement. We hypothesized an interaction effect between ethnicity and each predictor variable: self-efficacy, hope, and optimism in predicting academic achievement. The self-efficacy, hope, and optimism variables were centered and uncentered for comparison purposes. In addition, the variance inflation factor was calculated to determine the degree of multicollinearity present in the data results.

Participants

High school students ($N = 209$) from an ethnically diverse, working-class public high school in the southeast United States participated in this study. The ethnic breakdown of the participants was 105 African-American and 104 Caucasian American.

Materials

Students reported their gender, ethnicity, age, year in school, academic achievement (1 = mostly A's to 6 = half C's or lower), and educational goals (1 = not important to 8 = very important). The academic hope, academic self-efficacy, and optimism variables represented a cognitive set of competence measures in predicting academic achievement among this diverse high school population.

The *Academic Hope Scale* (AHS) measured academic hope components and is a sub-scale from the Domain Specific Hope Scale-Revised which measures hope in life areas including social relationships, family life, physical health, psychological health, work, romantic relationships, leisure activities, and religions/spiritual life with moderately high score reliability of .89 and above (Campbell & Kwon, 2001).

The *Academic Self-Efficacy Scale* (ASES) measured student beliefs about how they react to different academic tasks to succeed in academic achievement. The ASES is an excerpt from Bandura's Multidimensional Self-Efficacy Scale and has moderately high score reliability ranging from .69 to .85 (Zimmerman, Bandura, & Martinez-Pons, 1992).

The *Life Orientation Test* (LOT) measured dispositional optimism, or one's expectancies that he or she will experience positive outcomes with score reliability of .76 (Scheier & Carver, 1985).

Results

A multiple regression analysis was conducted using SPSS 16.0 with centered and uncentered variables. The interaction effects along with the variance inflation factors are reported separately.

Hypothesis 1: Is there a statistically significant interaction between academic hope and ethnicity in predicting academic achievement?

Hypothesis 2: Is there a statistically significant interaction between academic self efficacy and ethnicity in predicting academic achievement?

Hypothesis 3: Is there a statistically significant interaction between optimism and ethnicity in predicting academic achievement?

The results of all three hypotheses clearly indicated the importance of centering variables when including an interaction term, as noted by Aiken and West (1991). More importantly, the variance inflation factor was also affected if variables were not centered; thus falsely indicating multicollinearity.

Interpretation of Interaction Effects

Interpretation of results would be erroneous for interaction effects with uncentered variables. Centering reduces VIF to acceptable levels with academic hope interaction significant, but academic self-efficacy and optimism interaction not significant. We provide separate regression results for each ethnic group for interpretation of the interaction effects using centered variables. To compute the separate regression models, academic hope, self-efficacy and optimism scores were centered on the means for each of the ethnic groups. To compare and interpret the results of the separate regression models, an F-value was computed to test for differences between the separate regression models, and a t-value was computed to test for differences between the separate regression coefficients (Kleinbaum & Kupper, 1978).

The results in Table 4 demonstrate that separate regression models provide clarity of interpretation that single regression models with interaction term may not provide. First, although the results of Academic Hope confirm the interaction effect between ethnicity and Academic Hope (Table 1), the separate regression models highlight a magnitude of the interaction effect not present in the single regression model. Namely, Academic Hope explains nine times more variance in the Caucasian American student population than the African-American student population. Second, the difference in variance between ethnic groups is not examined in the single model with the interaction term, and as such may introduce Type II error. Although the interaction terms were not significant for Academic Self-Efficacy (Table 2) or Optimism (Table 3) in the single regression models, testing for differences in the amount of variance accounted for by each of the separate regression models reveals significant differences between Caucasian American and African-American students. Therefore, significant differences existed for academic self-efficacy and optimism not captured by the interaction effect in the single regression equation.

Summary and Conclusion

Historically, interaction effects and main effects were conducted using analysis of variance. In this context, a summary table reported the F values with a significant interaction effect being plotted to visually display an ordinal or disordinal interaction amongst the cell means. Multiple regression can also analyze various types of interaction effects, but how interaction effects are computed is important. We strongly recommend that variables be centered and the variance inflation factor reported otherwise erroneous results could occur and be misinterpreted.

Table 1. *Uncentered and Centered Regression Models of Academic Hope and Ethnicity Predicting Grade Point Average*

	<i>B</i>	<i>SE B</i>	<i>B</i>	<i>t</i>	<i>p</i>	<i>VIF</i>
Uncentered Regression Model						
Intercept	0.804	0.351		2.291	0.023	
Ethnicity	1.266	0.582	0.844	2.175	0.031	37.064
Academic Hope Scale (AHS)	0.057	0.009	0.507	6.168	0.000	1.668
Ethnicity x AHS	-0.036	0.015	-0.978	-2.453	0.015	39.198
Centered Regression Model						
Intercept	2.988	0.068		43.633	0.000	
Ethnicity	-0.134	0.097	-0.089	-1.379	0.170	1.027
Academic Hope Scale (AHS)	0.057	0.009	0.507	6.168	0.000	1.668
Ethnicity x AHS	-0.036	0.015	-0.201	-2.453	0.015	1.653

Note: $R^2 = 0.168$; $F(3, 205) = 13.802$, $p = 0.000$ for both uncentered and centered regression models

Table 2. *Uncentered and Centered Regression Models of Academic Self-Efficacy and Ethnicity Predicting Grade Point Average*

	<i>B</i>	<i>SE B</i>	β	<i>t</i>	<i>p</i>	<i>VIF</i>
Uncentered Regression Model						
Intercept	1.135	0.368		3.083	0.002	
Ethnicity	0.672	0.558	0.447	1.204	0.230	32.654
Academic Self Efficacy (ASE)	0.031	0.006	0.436	4.956	0.000	1.831
Ethnicity x ASE	-0.013	0.009	-0.535	-1.403	0.162	34.380
Centered Regression Model						
Intercept	2.957	0.069		42.590	0.000	
Ethnicity	-0.098	0.098	-0.065	-0.998	0.320	1.008
Academic Self Efficacy (ASE)	0.031	0.006	0.436	4.956	0.000	1.831
Ethnicity x ASE	-0.013	0.009	-0.123	-1.403	0.162	1.825

Note: $R^2 = 0.133$; $F(3, 205) = 10.51$, $p = 0.000$ for both uncentered and centered regression models

Table 3. *Uncentered and Centered Regression Models of Optimism and Ethnicity Predicting Grade Point Average*

	<i>B</i>	<i>SE B</i>	β	<i>t</i>	<i>p</i>	<i>VIF</i>
Uncentered Regression Model						
Intercept	1.758	0.383		4.591	0.000	
Ethnicity	0.223	0.538	0.149	0.415	0.679	28.346
Optimism (Opt)	0.044	0.014	0.300	3.108	0.002	2.061
Ethnicity x Opt	-0.011	0.020	-0.201	-0.547	0.585	29.795
Centered Regression Model						
Intercept	2.936	0.072		40.972	0.000	
Ethnicity	-0.066	0.101	-0.044	-0.654	0.514	1.001
Optimism (Opt)	0.044	0.014	0.300	3.108	0.002	2.061
Ethnicity x Opt	-0.011	0.020	-0.053	-0.547	0.585	2.059

Note: $R^2 = 0.071$; $F(3, 205) = 5.247$, $p = 0.002$ for both uncentered and centered regression models

Although multiple regression can analyze interaction effects, our results demonstrate that interaction terms not significant in regression models can produce significantly different regression models when computed separately. Therefore, we also suggest that separate regression equations be computed for each level of the interaction variable to provide a more robust interpretation of the interaction effect. This is a different research question than testing for an interaction effect, but traditional research methods dictate that non-significant interaction does not warrant further exploration, however, our results suggest otherwise.

Table 4. Summary of Separate Regression Analyses Predicting Grade Point Average

	African-American ($n = 105$)					Caucasian American ($n = 104$)				
	b	$SE\ b$	β	t	p	b	$SE\ b$	β	t	p
Academic Hope										
Intercept	2.855	0.069		41.557	0.000	2.988	0.068		43.725	0.000
Academic Hope Scale	0.020	0.012	0.170	1.750	0.083	0.057	0.009	0.522	6.181	0.000
Note: $R^2 = 0.029$; $F(1,103) = 3.1$, $p = 0.08$ for African-American students; $R^2 = 0.272$; $F(1,102) = 38.2$, $p = 0.0001$ for Caucasian American students; Comparison of regression models, $F(1, 207) = 68.57$, $p < 0.001$; Comparison of regression coefficients, $t = 3.7$, $p < 0.001$										
Academic Self-Efficacy										
Intercept	2.859	0.067		42.930	0.000	2.957	.072		41.113	0.000
Academic Self-Efficacy	0.018	0.007	0.256	2.687	0.008	0.031	.006	0.428	4.784	0.000
Note: $R^2 = 0.065$; $F(1,103) = 7.22$, $p = 0.008$ for African-American students; $R^2 = 0.183$; $F(1,102) = 22.89$, $p = 0.0001$ for Caucasian American students; Comparison of regression models, $F(1, 207) = 29.75$, $p < 0.001$; Comparison of regression coefficients, $t = 2.0$, $p < 0.05$										
Optimism										
Intercept	2.870	0.067		43.152	0.000	2.936	0.076		38.518	0.000
Optimism	0.033	0.013	0.247	2.586	0.011	0.044	0.015	0.278	2.922	0.004
Note: $R^2 = 0.048$; $F(1,103) = .24$, $p = 0.63$ for African-American students; $R^2 = 0.124$; $F(1,102) = 1.58$, $p = 0.21$ for Caucasian American students; Comparison of regression models, $F(1, 207) = 17.87$, $p < 0.001$; Comparison of regression coefficients, $t = .76$, $p > 0.10$										

References

- Aiken, L.S. & West, S.G. (1991). *Multiple Regression: Testing and Interpreting Interactions*. Sage Publications: Thousand Oaks, CA.
- Bandura, A. (1997). *Self-efficacy: The Exercise of Self-Control*. New York: Freeman.
- Belsley, D. A., Kuh, E. & Welsch, R. E. (1980). *Regression Diagnostics: Identifying Influential Data and Sources of Collinearity*. New York: John Wiley.
- Berry, W. D. & Feldman, S. (1985). *Multiple Regression in Practice*. London: Sage Publications.
- Campbell, D.G. & Kwon, P. (2001). Domain-specific hope and personal style: Toward an integrative understanding of dysphoria. *Journal of Social and Clinical Psychology*, 20(4), 498-520.
- Fox, J. (1997). *Applied Regression Analysis, Linear Models, and Related Methods*. Sage Publications: Thousand Oaks, CA.
- Freund, R.J., Littell, R.C., and Creighton, L. (2003). *Regression Using JMP*. Cary, NC: SAS Institute, Inc.
- Hinkle, D.E., Wiersma, W., & Jurs, S.G. (1998). *Applied Statistics for the Behavioral Sciences, 4th Edition*. Houghton Mifflin Company: NY.
- Kleinbaum, D. G., & Kupper, L. L. (1978). *Applied Regression Analysis and Other Multivariable Methods*. North Scituate, MA: Duxbury Press.
- Scheier, M.F., & Carver, C.S. (1985). Optimism, coping, and health: Assessment and implications of generalized outcome expectancies. *Health Psychology*, 4, 219-247.

- Scheier, M.F., & Carver, C.S. (1992). Effects of optimism or psychological and physical well-being: Theoretical overview and empirical update. *Cognitive Therapy and Research*, 16, 201-228
- Schroeder, L. D., Sjoquist, D. L. & Stephan, P. E. (1986). *Understanding Regression Analysis*. Beverly Hills, CA: Sage Publications.
- Schumacker, R.E. & Marcoulides, G.A. (1998). *Interaction and Nonlinear Effects in Structural Equation Modeling*. Erlbaum Associates: Mahway, NJ.
- Snyder, C. R. (1995). Conceptualizing, Measuring, and Nurturing Hope. *Journal of Counseling & Development*, 73, 355-360.
- Stine, Robert, A. (1995, February). Graphical interpretation of variance inflation factors. *The American Statistician*, 49(1), 53-56.
- Zimmerman, B., Bandura, A., & Martinez-Pons, M. (1992). Self-motivation for academic attainment: The role of self-efficacy beliefs and personal goal setting. *American Educational Research Journal*, 29, 663-676.

Send correspondence to: Randall E. Schumacker
 University of Alabama
 Email: rschumacker@ua.edu

Using Linear Regression and Propensity Score Matching to Estimate the Effect of Coaching on the SAT

Ben Domingue

Derek C. Briggs

University of Colorado

Using observational data from the Education Longitudinal Survey of 2002, the effect of coaching on the SAT is estimated via linear regression and propensity score matching approaches. The key features of taking a propensity score matching approach to support causal inferences are highlighted relative to the more traditional linear regression approach. A central difference is that propensity score matching restricts the sample from which effects are estimated to coached and uncoached students that are considered comparable. For those students that have taken both the PSAT and SAT, effect estimates of roughly 11 to 15 points on the math section and 6 to 9 points on the verbal are found. Only the math effects are statistically significant. We found that coaching is more effective for certain kinds of students, particularly those who have taken challenging academic coursework and come from high socioeconomic backgrounds. In the present empirical context the summary causal inference being drawn does not depend much upon whether the effect is estimated using linear regression or propensity score matching.

The SAT plays a high stakes role in the American college admissions process. As a consequence, a commercial industry has emerged to prepare students to take the SAT. According to The College Board, the SAT “tests students’ knowledge of subjects that are necessary for college success” (College Board, 2009). If short-term preparation (i.e., “coaching”) can be linked to substantial score increases, then the claim that the SAT measures knowledge that is developed gradually over years of schooling becomes equivocal. Moreover, if coaching services, some of which are quite expensive, give those who can afford them a sizable boost in their SAT scores, then this would further tilt an already uneven playing field when it comes to college admissions. One purpose of this study is to gauge the size of the coaching effect using data from a recent national cohort of high school students.

From an experimental design perspective, the most accurate method of estimating the causal effect of coaching would be through the use of a randomized experiment. In principle, randomization allows for relatively straightforward estimation of causal effects since, when done correctly, the process of randomization ensures that differences between experimental treatment and control groups on some outcome of interest can be attributed to the treatment, and not to some other variable. In short, a randomized experiment controls for confounding by design. Of course, randomized experiments are both expensive and hard to implement on a large scale. In this study we make use of observational data taken from the Educational Longitudinal Survey of 2002 (ELS:02). The ELS:02 data is strictly observational in the sense that students self-select to participate in test preparation programs. This presents substantial complications in drawing causal inferences since score differences between students who do and do not receive coaching may be confounded by preexisting differences on variables correlated with both coaching status and SAT performance. To account this we use linear regression (LR) and propensity score matching (PSM) in an attempt to control statistically for observable confounders in the process of estimating causal effects. A second purpose of this study is to compare and contrast the use of these two methods to support causal inferences in an observational setting.

We begin by describing the existing research on SAT coaching and highlight some methodological differences and similarities between LR and PSM. In the next section, we describe the ELS:02 data, presenting descriptive statistics for coached and uncoached students. Since propensity score matching is an umbrella term that encompasses a variety of different analytical procedures, we present the specifics of the two PSM approaches we will be invoking. Sections that focus on our empirical results, and the sensitivity of these results to our modeling assumptions follow. Finally, we compare this research to other studies in the SAT coaching literature and make a few methodological recommendations regarding the use of PSM to support causal inferences in education research contexts.

Background

Since 1953 there have been more than 30 studies conducted to evaluate the effect of coaching on specific sections of the SAT (Briggs, 2002). While one might assume from this that the empirical effectiveness of coaching on SAT performance has been well-established, this is only somewhat true. One principal reason for this is that the vast majority of coaching studies conducted over the 40 year period

between 1951 and 1991 tended to involve small samples that were not necessarily representative of the national population of high school seniors taking college admissions exams, and of the programs offering test coaching. In addition, a good number of these studies contained a variety of methodological flaws that compromised the validity of their conclusions. To date, the findings from analyses that are at both methodologically sophisticated and generalizable have indicated that the effect of coaching is about 10-20 points on the math section of the SAT and 5-10 points on the verbal section of the exam (Powers & Rock, 1999; Briggs, 2001; Briggs, 2002).

There are two principal motivations for conducting a new study to estimate the effect of SAT coaching using data from ELS:02. First, the SAT has undergone substantial changes to its format over the years (Lawrence et al., 2004). The SAT administered to students in the ELS:02 sample in 2003-2004 was considerably different in format than the SAT administered from 1953 through 1992 (the period during which the bulk of coaching studies were conducted). In particular, the types of items thought to be most coachable in the critical reading and math sections have been replaced (antonyms and analogies in the critical reading section; quantitative comparisons in the mathematics section). Because of this it is conceivable that the SAT has become less coachable over time, and this is one hypothesis we are testing in the present study.

Second, a relatively new method of estimating causal effects from observational data—propensity score matching (PSM)—has become increasingly popular over the past decade. The studies of Briggs (2001) and Powers and Rock (1999) both illustrate the classic approach of drawing inferences from observational data using a linear regression model (although both studies did use other methods as well): A single dummy variable represents treatment status and is included in a regression alongside other variables thought to be confounders. The estimated coefficient on the treatment variable represents the causal effect of coaching. There are, however, some clear limitations to the estimation of causal effects using linear regression. The biggest is that the method generally assumes that all potentially confounding variables have been measured without error and properly included in the model's specification (for details on the assumptions of the linear regression model in the context of estimating a coaching effect, see Briggs, 2004). Beyond this, the approach makes strong parametric assumptions. When many covariates are included in the model, the analyst is often implicitly relying upon assumptions of linearity and extrapolations well beyond observed variable combinations to correct for differences between treated and control units.

Rosenbaum and Rubin (1983) originally developed the idea and theoretical justification for PSM. The method tends to be consistent with an underlying framework for causal inference described by Holland (1986) as "Rubin's Causal Model." Recently the use of PSM within the framework of Rubin's Causal Model has become more visible in the education research literature. Hong and Raudenbush (2005, 2006) used PSM to evaluate the effect of kindergarten retention policies on academic achievement. Morgan (2001) used PSM to analyze the effect of Catholic schools on learning. Ruhm and Waldfogel (2007) used ratio matching techniques to estimate the effect of prekindergarten on later performance. A complete literature review would be beyond the scope of this paper, but these articles offer some indication of the range of educational questions that have been addressed using PSM techniques. Details for the particular analysis used here are given in the next section; a more general survey of PSM techniques is given in Caliendo and Kopeinig (2008). While PSM methods still make the assumption of "selection on observables," they typically relax the parametric assumptions associated with regression-based techniques, and perhaps most importantly, they focus the researcher's attention on the comparability of treatment and control units. Some subjects receiving an experimental treatment are simply not comparable to subjects receiving a control and vice-versa. Under a PSM approach subjects that are not comparable are excluded from the analysis and not used to estimate a causal effect.

The PSM approach has prompted strong claims from its proponents: "With flexible matching routines increasingly available, will regression adjustment for observational studies soon be obsolete?" (Hansen, 2004, p. 617). In fact, Hansen raised this question after performing a new analysis using the data from the Powers and Rock (1999) study on coaching effectiveness. Powers and Rock had used PSM methodology alongside regression to estimate the effect of SAT coaching, and the two methods had yielded similar results. In contrast, Hansen's estimates were roughly 5-8 points higher in math and 6 points lower in verbal, and he concluded that they were more defensible than those found by Powers and Rock¹. In this paper, we revisit this conclusion in the context of comparing the LR and PSM approaches to estimating the effect of coaching on SAT performance.

ELS Variables

Our analysis is based upon data from the ELS:02 survey conducted by the National Center for Educational Statistics. The survey followed a longitudinal cohort of high school sophomores in 2002 through their senior years (2004) and beyond (2006). It is designed to be representative of this national cohort of American high school students. The ELS:02 data contains basic demographic information about students, as well as more specific information about their academic achievement, attitudes, and opinions on a variety of subjects related to their schooling experiences. Information in the dataset on students' grades, test scores, and course-taking are based on official high school transcripts and test reports from the test makers rather than being self-reported.

The "treatment" of interest in this study is defined by the responses to a set of questions that asked students whether and how they prepared for the SAT. Students were able to indicate if they had prepared through the use of school courses, commercial courses, tutoring, or a variety of preparatory materials. In what follows, a coached student is defined as one that reported participating in a commercial preparatory course. Of the 16,197 students in the ELS:02 data, we restricted the sample to those students who had a 10th grade transcript available, responded to both the 2002 (grade 10) and 2004 (grade 12) surveys, and took the PSAT and SAT. We refer to these students as the "POP1" sample (N = 1,644). In contrast, we run separate analyses for a "POP2" sample (N = 2,549) that represents those students who took the SAT but did not take the PSAT². There are some distinct differences between the types of students who take the PSAT and those who do not. Students that take the PSAT are much more likely to be college-bound and motivated to perform well on the SAT. Splitting the sample into the POP1 and POP2 groupings allows for explicit comparison to the results from Briggs (2001), where the same groups of students were defined from an earlier survey of a longitudinal student cohort from 1988 to 1992 (NELS:88). However, from the perspective of drawing unbiased causal inferences, the POP1 sample is clearly preferable to the POP2 sample because PSAT scores (available for the former but not the latter) are well correlated with both coaching status and subsequent SAT performance.

Descriptive statistics for the variables used in this analysis are given in Table 1 (an index of the variables used in this study is given in the Appendix). What sorts of variables that may be correlated with SAT performance serve to distinguish students that do and do not participate in commercial coaching? As elaborated in previous work (Briggs, 2002) these sorts of variables fall into roughly three groups: demographic characteristics of students, variables that proxy for academic achievement and aptitude, and motivational variables. In Table 1 we can see to what extent coached and uncoached students differ with respect to these variables. In many cases the differences are significant. For example, relative to uncoached students, coached students in POP1 score 2.3 points better on the PSATM and 1.6 points better on the PSATV (or 23 and 16 points when expressed on the SAT scale shown in Table 1). For students in the POP2 sample such comparisons with respect to PSAT obviously cannot be made, however, students participating in the ELS base year survey (in grade 10) were administered standardized tests in both math (BYMATH) and reading (BYREAD). These tests have strong positive correlations with the PSAT, so for students in the POP2 sample they serve as a substitute. However, it seems clear that they are an imperfect substitute. For the POP1 sample, in contrast to the mean differences observed on PSAT scores, there is no significant difference in mean BYMATH and BYREAD scores between coached and uncoached students. Hence it is likely that the BYMATH and BYREAD variables used for the POP2 sample do not fully capture the differences in prior ability to perform well on high stakes tests that is captured by the PSAT variable.

Major differences also exist between these two groups in terms of socio-economic status (SES), GPA (especially in POP2), group percentage of Asian students, percentages in private and rural schools, percentages of remedial course takers, percentages of ESL students, and percentages taking college preparatory curriculum and doing more than 10 hours per week of homework. In contrast, coached and uncoached students have fairly similar numbers of math credits and attend urban schools in similar percentages (especially in POP1). Student motivation is classic example of a plausible confounding variable in the context of coaching studies. Since the SAT is viewed by students as a crucial piece of their college application, they may have far greater motivation to perform well on this test than previous

Table 1. Variable Means for POP1 and POP2 by Coaching Status

Variable & Brief Description	POP1 ^a			POP2 ^b		
	Coached	Uncoached	p value	Coached	Uncoached	p value
PSATM*10 ^c	542	519	0	NA	NA	NA
PSATV*10 ^c	524	508	0	NA	NA	NA
BYMATH-ELS Math Test	57.6	57.2	0.19	56.2	55.8	0.19
BYREAD-ELS Reading Test	57.2	56.9	0.25	55.2	55.1	0.46
AGE/12 ^c	17.8	17.9	0.26	17.9	17.8	0.10
SES Index	0.68	0.38	0	0.56	0.3	0
GPA	3.09	3.05	0.14	3.08	2.96	0
MCRD-# of math credits	3.84	3.81	0.28	3.77	3.76	0.44
Below variables are categorical, means expressed as percents.						
FEMALE	53	56	0.15	57	50	0
ASIAN	24	12	0	28	16	0
BLACK	12	8	0.01	15	14	0.36
NATIVE	3	3	0.48	4	5	0.18
HISPANIC	6	9	0.06	12	11	0.22
PRIVATE	56	43	0	36	25	0
RURAL	5	12	0	10	18	0
URBAN	39	39	0.50	42	33	0
AP-Taken an AP course	58	52	0.01	64	48	0
REM_ENG-Remedial English	7	6	0.33	6	7	0.21
REM_MATH-Remedial Math	8	7	0.18	6	8	0.17
COLL_PREP-HS curriculum	78	75	0.12	76	72	0.05
HW->10 hours/wk homework	40	25	0	34	22	0
ESL	23	13	0	23	16	0
EDU_AFTER_HS ^d	96	88	0	92	86	0
COLLEGE_INFO ^d	6	11	0	6	10	0
PRNTS_DISC_PREP ^d	34	19	0	37	22	0
PRNTS_DISC_SCH ^d	39	33	0.02	45	37	0
NERVES.M ^d	1	2	0.05	NA	NA	NA
NERVES.V ^d	2	2	0.34	NA	NA	NA
UNDERPERFORM.M ^d	14	14	0.44	13	16	0.03
UNDERPERFORM.V ^d	13	15	0.22	14	16	0.08
N	357	1195		448	1941	

a. All students who responded to 2002 and 2004 surveys, had 10th grade transcripts, and took the PSAT and SAT.

b. All students who responded to 2002 and 2004 surveys, had 10th grade transcripts, and took the SAT (but not the PSAT).

c. PSATM, PSATV, and AGE are all transformed for the table, but the untransformed variables were used in the analysis.

d. See Appendix.

standardized tests (such as the PSAT or the ELS base year tests) or in their classes (as expressed by their GPA). Variables such as whether or not a student plans to continue his/her education after high school (EDU_AFTER_HS) and whether or not a student has sought out information about college (COLLEGE_INFO) capture some aspects of student motivation. We also created two new variables as proxies for motivation: UNDERPERFORM and NERVES. Students with a low GPA (less than 3.0) but a math or verbal score greater than the mean (for POP1 or POP2) on the ELS base year test were given a 1 on the variable UNDERPERFORM. We reasoned that such students have greater academic ability than their GPA suggests, and may sense a greater need to perform well on the SAT as the deadlines for college admissions approach. Should these students elect to get coached in preparation for the SAT, part of any score increase may be due to their new-found motivation rather than coaching. Those students who did substantially worse on the PSAT than we would have predicted given their performance on the ELS base year tests get a value of “1” on the variable NERVES³. Such students, doing worse on the PSAT than they may have expected, may be quite likely to sign up for coaching. We may falsely attribute a later score gain on the SAT for these students to coaching when we are in fact merely seeing a regression to the mean. The second variable relied upon a student’s PSAT score, so we created it only for those students in POP1. As can be seen in Table 1, we found significant differences between coached and uncoached students on some of these variables as a function of SAT test subject and POP1 or POP2 membership.

One complication in using the ELS data is that there is a substantial amount of missing data. To simply exclude the missing cases would not only eliminate a large percentage of our data, but would also necessitate either the “missing at random” or “missing completely at random” assumptions (Rubin, 1976) which may be difficult to support. In examining the variables we found that certain dichotomous variables contain the bulk of the missing data⁴. Rather than simply throwing these cases out, we included missingness as an additional level of these variables⁵. This strategy, also followed by Hansen (2004), allowed us to include most of the POP1 and POP2 samples in our analysis. Only students with missing data on our continuous variables were removed from the subsequent analysis. This decreased the POP1 sample from 1,644 students to 1,552 and the POP2 sample from 2,549 to 2,389. While this is not the only approach available for dealing with missing data (imputation techniques would also be an option), this did allow us to retain most of our sample without a substantial increase in the complexity of our analysis. One final complication was that each student who failed to respond to the question regarding being Black also failed to respond to the question about being of Native American origin. Since this led to linearity between these variables, they were aggregated into a variable called RACE.

Method

Because the approach typically used in a regression analysis is well understood, we will focus here on the methods used in our PSM analysis. Similar to Caliendo & Kopeinig (2008), we break a PSM analysis into five separate steps:

1. Estimating the Propensity Score
2. Implementing a Matching Algorithm
3. Assessing the Balance after Matching
4. Computing an Effect Estimate
5. Sensitivity Analysis

We shall discuss these five steps as we implement them in our analysis that follows.

The propensity score is at the core of the PSM methodology. It is the estimated probability of the unit of analysis receiving the treatment given the observed covariates, typically computed using logistic regression. Unbiased estimation of causal effects relies upon selection into treatment being a function of only those covariates used in the estimation of propensity scores. What to include in the selection function, the function which predicts treatment status, and how to choose its functional form are aspects of the methodology about which there is still some confusion in the literature. However, one agreed upon aspect in the PSM literature is that the success of the selection function should be the “balance” it generates in the distribution of covariates among treatment and control groups that have been matched according to their propensity scores.

Researchers have used the estimated propensity scores to compare treatment and control groups in several ways—inverse propensity weighting and kernel matching being two alternatives (see Frank et al., 2008 for an example of the first and Callahan et al., 2009 for an example of the second)—but in the present students we focus on two matching approaches that appear commonly in education research

applications: subclassification and optimal pair matching. In the subclassification approach, units in the common support (the area of overlap on the estimated propensity score between the coached and uncoached groups) are split into subclasses based upon the quantiles of the distribution of the estimated propensity scores for the coached students. In the process we eliminate those students from the analysis who lack directly comparable counterparts in terms of their propensity scores. As noted earlier, this constitutes a key distinction of PSM relative to LR. (As it turns out, in this study there is good overlap between coached and uncoached students on the estimated propensity scores, so we do not lose substantial portions of the POP1 or POP2 groups.) The optimal pair match is a one-to-one matching algorithm, meaning that each coached student is matched to a single uncoached student. A consequence of pair matching is that a larger number of uncoached students will be excluded from the analysis relative to subclassification. Optimal matching is done such that we obtain the lowest possible mean difference across all of the matches, hence the use of the term “optimal”. We used the R statistical computing package (R Development Core Team, 2008) as well as specialized matching software (Ho et al., 2004) to perform the propensity score estimation and matching.

Propensity scores are a means to an end. They are used to match treatment and control units such that after the units have been matched, their covariate distributions will be equivalent. For example, prior to matching, coached students may have higher mean PSAT scores than uncoached students. After matching, the PSAT means (and SDs) for coached and uncoached students should be about the same. Balance is a necessary condition for unbiased estimation via PSM. We apply two approaches to evaluate balance: the reduction in standardized mean differences after matching and an omnibus test for balance (Hansen & Bowers, 2008). Standardized differences are differences in coached versus uncoached covariate means relative to a weighted combination of the standard deviations across matched units. The omnibus test for balance is designed to answer the following question. Is the degree of difference between the treatment and controls groups consistent with that which would be expected between two groups randomly created from a single sample, as in an experiment? The test statistic we use, computed via the software of Bowers et al. (2008), simultaneously tests a Fisher Randomization Hypothesis for all covariates.

If the matched data looks as though it could have come from randomization, the simplest approach to computing an effect estimate is to compare differences in group means. If selection into treatment is solely a function of the observable data, then this will be an unbiased estimate of the treatment effect (Rosenbaum & Rubin, 1983). This is the basic idea behind the approach we use to estimate the coaching effect under the subclassification PSM approach. In contrast, for the optimally pair matched data, we depart from this simple approach because clear imbalances remain on important covariates (the PSAT for the POP1 sample and the ELS base year tests for the POP2 sample), even after matching. Hence we estimate the coaching effect after adjusting for remaining PSAT score differences using a regression model. We estimate standard errors using the Huber-White correction to adjust for the clustering of students at the school level.

As with the LR approach, the validity of the PSM-based estimates for the coaching effect depend most fundamentally upon the availability of all the relevant variables that predict whether or not a student is likely to be coached. Rosenbaum (2002) outlines a procedure which allows us to check the robustness of our results to certain deviations from this assumption. In particular, we assume that there exists a hidden variable with a known relationship to treatment status. Using Keele’s (2008) software, we are able to examine the degree to which our effect estimates may change if such a confounder were to exist.

Results

Linear Regression

In presenting the results, we first discuss the results from the regression analysis. Table 2 shows the coefficient estimates for each of four regressions: both sections of the test for both POP1 and POP2 samples. The estimated effect of coaching on the math section is 11 and 22 points for the POP1 and POP2 samples respectively, both statistically significant at the .01 level. For the verbal section, the effects were 6 points for POP1 and 8 points for POP2, only the second of which was statistically significant at the .05 level.

Table 2. Coefficient Estimates with Huber-White Standard Errors for Linear Regression Analysis

Variable	SATM				SATV			
	POP1		POP2		POP1		POP2	
	Coeff	SE	Coeff	SE	Coeff	SE	Coeff	SE
(Intercept)	41.7	51.7	63.0	55.5	48.5	45.4	29.8	61.9
COACH	11.3	3.1	21.5	3.0	5.8	3.6	7.9	3.8
PSATM	5.0	0.3	NA	NA	0.5	0.3	NA	NA
PSATV	1.0	0.3	NA	NA	6.4	0.3	NA	NA
BYMATH	3.3	0.4	7.6	0.3	0.8	0.3	2.3	0.3
BYREAD	0.0	0.3	0.6	0.2	1.8	0.3	5.3	0.3
AGE	-0.2	0.3	-0.4	0.2	-0.2	0.2	-0.1	0.3
SES	1.4	2.3	10.2	2.1	7.1	2.1	18.1	2.4
FEMALE	-13.0	2.7	-18.0	3.0	0.2	2.9	-6.5	2.9
ASIAN1	0.9	5.6	10.5	4.7	-2.2	5.4	-13.4	4.5
ASIANNA	-41.6	31.6	-71.0	70.9	19.4	27.9	-57.2	49.8
RACENATIVE	6.8	6.7	-0.2	5.9	1.8	10.1	-0.3	6.7
RACEBLACK	-11.4	4.0	-16.1	4.9	8.2	5.4	-4.5	5.2
RACEBOTH	-28.9	11.6	-9.5	13.5	0.4	12.3	-4.7	10.7
RACEBOTHNA	31.3	31.2	72.9	70.4	-21.8	26.7	52.2	49.4
HISPANIC1	-9.8	4.3	-17.0	4.9	0.5	7.4	-2.9	6.5
HISPANICNA	2.3	18.2	12.4	13.2	21.2	16.4	-17.2	9.4
PRIVATE	1.6	3.2	8.4	3.7	2.7	3.5	16.5	4.3
RURAL	-4.7	4.9	-2.1	3.9	-4.3	4.2	2.5	4.6
URBAN	-6.5	3.4	-1.0	3.5	-10.4	3.7	-2.4	3.9
AP	12.5	3.2	27.0	3.2	12.6	3.4	33.7	3.5
REM_ENG1	12.1	9.1	17.3	8.0	2.7	9.3	-3.5	8.7
REM_ENGNA	-5.7	14.4	2.4	12.2	7.7	13.3	-10.3	13.6
REM_MATH1	-12.4	9.0	-23.6	7.6	-0.9	7.2	7.6	9.1
REM_MATHNA	11.9	15.8	5.1	12.3	-1.3	12.8	16.4	12.7
COLL_PREP	-1.0	3.4	-0.6	2.8	-1.4	3.2	6.7	3.0
MATH_CRD	0.5	1.8	2.8	1.4	0.4	1.8	-2.2	1.5
HW	0.6	3.3	8.5	2.8	2.2	3.0	3.7	3.2
ESL	17.8	4.3	8.0	4.0	2.3	4.5	-8.6	4.9
EDU_AFTER_HS	1.5	4.2	0.5	3.5	-0.4	4.4	-6.5	4.1
COLLEGE_INFO1	-4.4	5.8	1.8	4.5	3.1	4.4	4.0	4.9
COLLEGE_INFONA	4.5	6.6	16.5	7.3	-0.1	8.1	12.8	7.1
PRNTS_DISC_PREP1	1.9	3.8	5.8	3.4	5.3	3.6	-0.8	3.8
PRNTS_DISC_PREPNA	-9.7	9.2	-21.0	10.0	13.0	9.6	-6.0	11.8
PRNTS_DISC_SCH1	-7.6	3.0	-8.1	2.8	-5.9	3.1	-3.4	3.1
PRNTS_DISC_SCHNA	16.3	10.1	18.4	10.8	-19.2	9.8	-8.2	12.4
GPA	14.4	3.8	23.3	3.4	7.9	3.5	24.6	3.4
UNDERPERFORM.M	3.0	5.7	9.5	4.6	1.9	4.7	1.3	5.4
UNDERPERFORM.V	-2.1	4.6	-10.5	4.2	-3.2	4.3	-10.0	4.9
NERVES.M	25.9	13.3	NA	NA	7.9	11.8	NA	NA
NERVES.V	12.5	10.2	NA	NA	25.0	11.4	NA	NA
N	1552		2389		1552		2389	
R ²	0.77		0.73		0.77		0.67	

Since the SAT scale has been internalized by many who have been educated in the United States, the magnitude of our causal effect has inherent meaning to those who recall taking the test as a high school student. However, there are two additional ways of contextualizing this effect. We can first compare our effect estimates to the unadjusted difference in SAT means for coached and uncoached students to assess the impact of adjusting for confounding variables. On the math section, there were unadjusted differences of 31 and 36 points in POP1 and POP2 respectively. After adjustment using LR, those differences fall to 11 and 22 points respectively. On the verbal section, initial differences of 23 and 22 points were reduced to 6 and 8 points for POP1 and POP2. Aggregating the effects across both parts of the test, we find that of the initial 54 point difference between coached and uncoached students in POP1, only 17 points can be attributed to coaching. For POP2, we can only attribute 30 of the 58 point initial difference to coaching. We can also express the coaching effects as effect sizes. Using the standard deviation of the uncoached students, the effect sizes of coaching on the math part were 0.11 and 0.20 for POP1 and POP2 samples. On the verbal section, the effect sizes were .06 and .07 for POP1 and POP2 samples respectively.

Propensity Score Matching

As a first step in our PSM analyses, we ran a logistic regression using the full set of covariates shown in Table 2. Logistic regression coefficients for the subclassification and pair matching approaches are shown in Table 3. We distinguish between the logistic regression coefficients for subclassification and optimal pair matching because in each case, different samples of students were included or excluded in the matching procedure depending upon where they fell within the area of common support (or overlap) on the estimated propensity score, and whether (in the case of the optimal matching approach) an uncoached student could be matched to a specific coached student. In the optimal pair matching approach, we only excluded those coached students from matching who were more likely to be coached than any uncoached student. In the subclassification approach, we also excluded students who were more likely to be uncoached than any coached student. Once students were excluded, the logistic regression was run again using the restricted sample. Matches in the optimal pair match were made using logit units rather than the actual propensity scores. The difference between similar propensity scores on the high or low end of the propensity score scale (near 0 or 1) will increase when logit units are used instead. Since the goal is to have the smallest global difference, using the logit forces better matches at the high and low end of the propensity score scale. Following the lead of Rosenbaum & Rubin (1983), we use quintiles of the propensity score distribution to form our subclasses⁶.

The focus in PSM on comparing only those units who are directly comparable on the estimated propensity score is important, so we draw some attention to the students who were “unmatchable” and hence, unlike in LR, were excluded as a basis for the subsequent estimation of a causal effect. Under subclassification, both coached and uncoached students who fall outside the region of common support were excluded. In POP1, this led to the exclusion of 54 uncoached students and 4 coached students. The coverage of the common support was even better in POP2, only necessitating the exclusion of 4 students, of which only one was coached. In the optimal pair match we only removed 4 coached students from POP1 and 1 coached student from POP2. After matching, 842 of the uncoached students from POP1 were not matched to coached counterparts, and 1,494 uncoached students from POP2 were not matched, an illustration of the restrictive nature of the pair matching approach.

As a first empirical evaluation of the extent to which balance has been achieved, Figure 1 shows plotted density curves of propensity score distributions for both POP samples under each matching algorithm. As we would expect, the top row of figures indicates that there are higher percentages of coached students who were likely to be coached (i.e., had higher propensity scores) on the basis of our predictor variables. The relative paucity of uncoached students in the higher range of the estimated propensity score indicates that the causal estimates for these subclasses, especially the highest subclass (notice that the vertical lines represent the subclass divisions), will rest upon comparisons of many coached students to few uncoached students.

Balance requires more than similar distributions for the estimated propensity scores among coached and uncoached groups; it is also necessary for the distributions of each relevant covariate to be similar. Figures 2 and 3 illustrate graphically the improvement in standardized mean differences after matching. In all four cases, there are generally big improvements in balance after matching. Balance appears to be strongest under the subclassification approach, where differences in covariates between the groups are

Table 3. Logistic Regression Coefficients for Subclassification

Variable	POP1		POP2	
	Coeff	SE	Coeff	SE
(Intercept)	-2.84	3.01	-5.25	2.38
PSATM	0.02	0.01	NA	NA
PSATV	0.00	0.01	NA	NA
BYMATH	-0.03	0.02	-0.01	0.01
BYREAD	0.00	0.01	-0.02	0.01
AGE	0.00	0.01	0.02	0.01
SES	0.59	0.11	0.50	0.09
FEMALE	-0.07	0.15	0.20	0.12
ASIAN1	0.78	0.24	0.70	0.17
ASIANNA	0.11	1.37	-0.20	0.32
RACENATIVE	-0.56	0.51	-0.11	0.30
RACEBLACK	0.60	0.24	0.45	0.18
RACEBOTH	2.50	0.81	-0.17	0.64
RACEBOTHNA	0.02	1.30	NA	NA
HISPANIC1	-0.32	0.33	0.43	0.22
HISPANICNA	-0.13	0.83	0.19	0.53
PRIVATE	0.56	0.16	0.40	0.14
RURAL	-0.65	0.30	-0.18	0.19
URBAN	-0.33	0.15	0.05	0.13
AP	0.01	0.17	0.54	0.14
REM_ENG1	-0.16	0.46	-0.28	0.39
REM_ENGNA	-0.23	0.90	-0.41	0.66
REM_MATH1	0.34	0.42	-0.03	0.37
REM_MATHNA	0.63	0.89	-0.20	0.65
COLL_PREP	-0.04	0.17	0.05	0.13
MATH_CRD	0.04	0.09	-0.04	0.06
HW	0.46	0.15	0.25	0.13
ESL	0.48	0.23	0.16	0.18
EDU_AFTER_HS	1.10	0.33	0.29	0.20
COLLEGE_INFO1	-0.28	0.27	-0.42	0.24
COLLEGE_INFONA	-0.48	0.40	0.53	0.30
PRNTS_DISC_PREP1	0.62	0.17	0.49	0.13
PRNTS_DISC_PREPNA	-0.64	0.62	-0.11	0.45
PRNTS_DISC_SCH1	0.02	0.16	-0.05	0.13
PRNTS_DISC_SCHNA	1.29	0.62	-0.53	0.48
GPA	-0.10	0.16	0.10	0.14
UNDERPERFORM.M	0.12	0.25	-0.05	0.22
UNDERPERFORM.V	-0.22	0.25	0.19	0.22
NERVES.M	-0.12	0.68	NA	NA
NERVES.V	0.33	0.52	NA	NA
N	1494		2385	

consistently less than one tenth of an SD. Next we applied the omnibus test described by Hansen & Bowers to evaluate whether the degree of balance that we observe is close enough to that which would be expected from the creation of two samples from a single population via random assignment. The result in each case was a test statistic with a very low p-value, which indicates that the balance within subclasses was similar to what one would expect had coached and uncoached students been randomly assigned (this was true for matches with both POP1 and POP2 samples). Interestingly, however, while the omnibus test suggests that both approaches result adequate balance, under the optimal matching

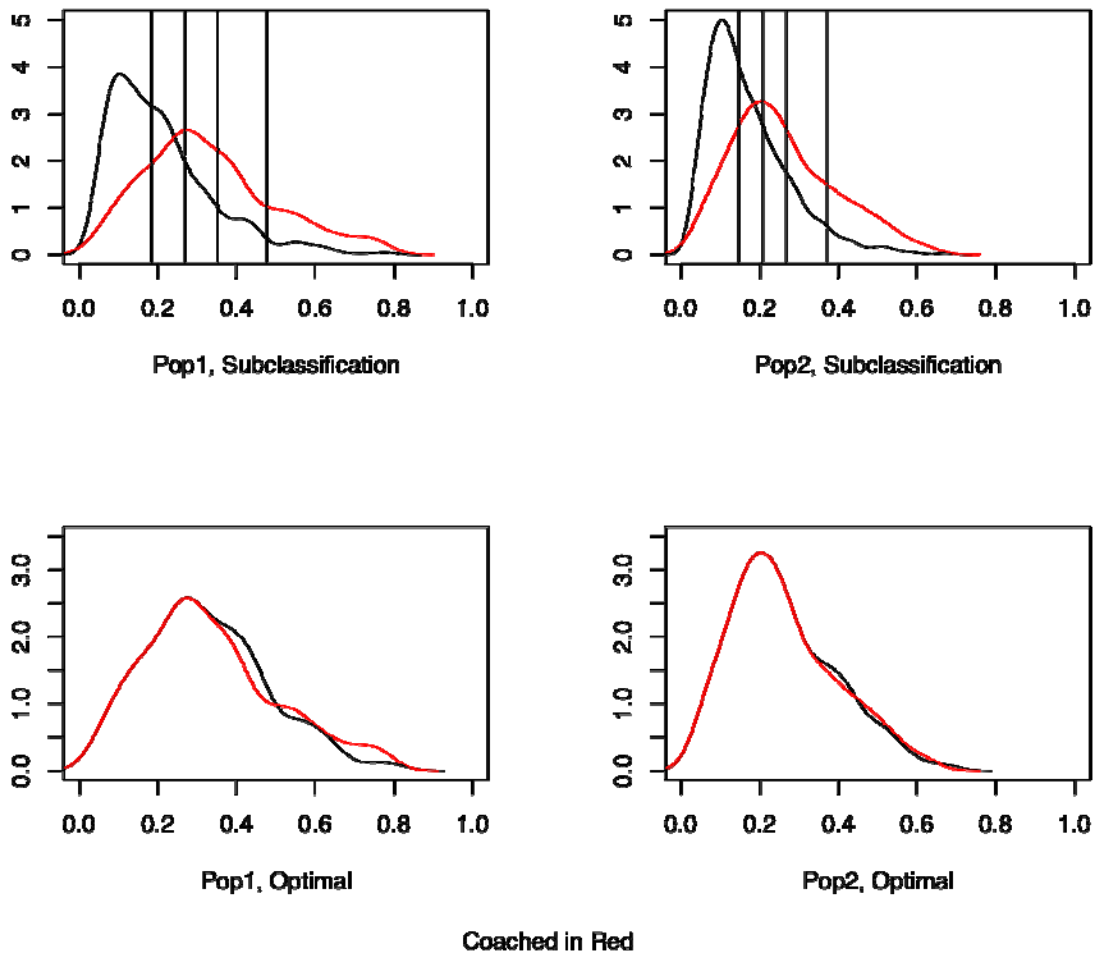


Figure 1. Density Curves for the 4 Matched Data Sets

approach there are clearly a number of variables for which important differences between the two groups remain. Most notably for the POP1 samples, coached students continue to have higher PSAT scores than their uncoached counterparts. For the POP2 samples, scores on the ELS base year tests in math and reading (BYMATH and BYREAD) actually become somewhat less balanced between coached and uncoached students after matching. These imbalances point to potential sources of bias that would still need to be taken into account when estimating coaching effects, and they also underscore the importance of not relying solely on tests of significance to support a conclusion that all plausible confounders are suitably balanced.

To compute effect estimates from our subclassified data we regressed the math or verbal portion of the SAT on a dummy variable indicating coaching status as well as a dummy variable indicating the propensity score subclass for the student. For math, the estimated effects were 12 and 22 points in POP1 and POP2 respectively. For verbal, they were 6 and 9 points. Only the POP2 effect for math was significant at the .05 level. In the optimally pair matched data, we began by regressing the sections of the test on a dummy variable for coaching status. This led to effects of 24 and 18 points in math for POP1 and POP2 and 14 and 0 points in verbal. However, due to the remaining imbalances on the PSAT tests shown in Figure 3, we ran subsequent regressions in which the variables PSATM and PSATV were included as controls for the POP1 sample, and the variables BYMATH and BYREAD were included as controls for the POP2 sample. In these regressions, the estimated math and verbal effects for POP1 fell from 24 and 14 points to 15 and 9 points. In contrast, because uncoached students had higher mean BYMATH and BYREAD scores than coached students after matching, the math and verbal effects increased from 18 and 0 points to 24 and 6 points after regression adjustment. These results are summarized in Table 4. Most of the coaching effect estimates for math are statistically significant at the .05 level in both POP1 and POP2 samples; in contrast, none of the verbal estimates meet this conventional threshold⁷.




Figure 2. Improvement in Standardized Differences after Subclassification




Figure 3. Improvement in Standardized Differences after Optimal Pair Match

One benefit of the subclassification approach is that it allows us to further probe the LR constraint that the effect of coaching can be best summarized with a single number. Table 5 shows the estimated coaching effect in each subclass as defined by ranges of the propensity score distribution. It also shows the ceiling for the estimated propensity scores in that subclass as well as the mean SES for the students in each subclass. Note that the effects vary quite dramatically with the

Table 4. Effect Estimates from Matching with Huber-White Standard Errors in Parentheses

Matching Method	Math		Verbal	
	POP1	POP2	POP1	POP2
Subclass	12.4 (8.4)	21.9 (6.5)	6.3 (8.5)	8.6 (6.6)
N	1494	2385	1494	2385
Optimal	23.5 (10.3)	18.0 (9.1)	14.2 (10.5)	0.3 (8.9)
Optimal with Smoothing	15.1 (4.7)	24.2 (4.5)	8.5 (4.4)	6.4 (5.7)
N	706	894	706	894

highest effect estimates found in the higher subclasses. The higher subclasses contain those students most likely to be coached. Since the propensity score correlates strongly and positively with SES, one plausible explanation is that more affluent students are buying coaching that is both more expensive and more effective.

Sensitivity Analysis

The coaching estimates from our original regression analyses were predicated upon the constraint that there is a single mean causal effect that applies to all students. As the results from the subclassification analysis above indicate, this may be unreasonably restrictive. Through the inclusion of interaction terms in the regression model, we can readily evaluate the possibility of differential causal effects of coaching for selected subsamples of students. In what follows we report these results (summarized in Tables 6 & 7) only for student subgroups in the POP1 sample. We focus on the POP1 sample because we are more confident in these coaching estimates since they control for preexisting differences in PSAT performance.

One characteristic of coaching services that our treatment variable cannot operationalize is that these programs vary widely in cost. Extremely expensive small group work with college professors may be classified as coaching right alongside much less expensive coaching offered in community centers. The effects of these two types of coaching may be quite different. Since we do not have information on how much money students paid for coaching, we instead use a dummy variable indicating whether a student was in the top quartile of the distribution for the SES variable supplied in the ELS:02 data. Those students that were in the top quartile of the SES index are potentially paying for much more expensive (and possibly higher quality) coaching. When coached students in the top SES quartile are compared to uncoached students in the top SES quartile, we find a mean difference of 15 points on SAT math scores. In contrast, amongst students in lower SES quartiles, coaching has only a 5 point effect. Hence math coaching appears to be 10 points more effective for high SES students than it is for lower SES students. On the verbal section the coaching by SES interaction is weaker: the effect is 9 points for high SES students but just 2 points for lower SES students.

There appear to be no gender related differences in the effectiveness of coaching. On the other hand, there are some important differences as a function of race/ethnicity. Coaching is differentially effective for Asian students, for whom the coaching effect is 5 and 14 points higher on the math and verbal sections respectively. One explanation for this is that 70% of the coached Asian students are also ESL students. For Black students, the effect of coaching on the math portion of the test was 15 points higher than it was for non-Black students. Conversely, the effect on the verbal section of the SAT for Black students was 18 points lower than the coaching effect for non-Black students. The latter result is driven by our finding of a negative effect for Black students on the verbal section of the SAT. Finally, students with AP course experience seem to benefit considerably from coaching. On each section of the SAT, we found a coaching effect that was 12 points higher than similar coached students who had never taken an AP course.

Table 5. Effect Estimates across Each Subclasses

Subclass	POP1				POP2			
	Class Ceiling	SES	Math	Verbal	Class Ceiling	SES	Math	Verbal
1	0.18	0.17	-18	-13	0.15	0.03	3	-20
2	0.27	0.55	9	-15	0.21	0.39	23	33
3	0.35	0.75	37	25	0.27	0.57	27	14
4	0.48	0.81	36	41	0.37	0.75	58	31
5	0.85	1.07	72	65	0.71	1.02	72	50

Turning now to a sensitivity analysis of our PSM results, we consider the effects of some alternate specifications and assumptions. A potential mistake one may make in a PSM analysis is to misspecify the selection function. We established the selection function used in our analyses through careful consideration of what variables should theoretically influence coaching status and SAT results. Once this has been established, in our view there is no sound rationale for excluding such variables from the selection function on the basis of their statistical significance. Nonetheless, such exclusions are possible and even likely whenever the variables for a selection function are chosen by a stepwise algorithm⁸. To assess how sensitive our results are to the exclusion of certain variables from the selection function on the basis of their statistical significance, we estimated coaching effects after matching on the basis of a restricted selection function. In this approach only those covariates which were significant at the .05 level in the initial logistic regression were retained and then the logistic regression was estimated again to generate new propensity score predictions. After matching using the same subclassification and pair matching techniques described above, we computed new causal effect estimates (summarized in Table 8). Most notably, the estimated coached effect was 5 points higher for math (from 12.4 to 17.4) under subclassification (POP1 sample), and 6 points higher for verbal (from 6.4 to 12.1) under optimal matching (POP2 sample). In all other cases the coaching effects were about the same using either selection function. While a 5 to 6 point change to the effect may seem relatively small as a “worst case scenario”, expressed as a percentage of the original effect (where the “original” effect is that deriving from the unrestricted selection function) they represent substantial increases of 40% and 89% respectively. Indeed, the differences that arise from this adjustment to the selection function are about as big as the largest differences found when shifting from an LR to a PSM approach to estimate coaching effects.

Another type of sensitivity analysis is to ask how our effect estimates would change in the presence of hidden bias

Table 6. Interaction Effects on SATM for Selected Variables (POP1 Sample)

Variable	Effect Outside ¹	Effect Inside ²	Difference
HI_SES	4.7	15.1	10.4
AP	6.5	14.9	8.3
FEMALE	11.3	11.2	-0.1
ESL	10.6	14.1	3.5
ASIAN	11.1	16.2	5.1
BLACK	11.4	26.4	15.0

Table 7. Interaction Effects on SATV for Selected Variables (POP1 Sample)

Variable	Effect Outside ¹	Effect Inside ²	Difference
HI_SES	1.8	9.4	7.6
AP	3.7	7.5	3.8
FEMALE	5.7	6.0	0.3
ESL	4.8	10.4	5.7
ASIAN	2.6	17.1	14.5
BLACK	8.5	-9.4	-18.0

¹ These columns show the coaching effect considering only those students outside the group of interest.

² These columns show the coaching effect considering only those students inside the group of interest.

(Rosenbaum, 2002). This method can be used to estimate the upper and lower bounds for how our effect estimate under the pair matching approach would change if we did not observe a variable which was predictive of treatment status. While the technical details of how such an analysis is conducted are outside the scope of this paper, we have found that our results are indeed sensitive to having fully observed all relevant confounders. In the presence of a

moderate hidden bias⁹, the effect of coaching in the POP1 sample may range anywhere between 5 to 45 points on the math section of the SAT and 0 to 25 points on the verbal section.

Selection Function	SATM		SATV	
	POP1	POP2	POP1	POP2
	Subclassification			
Original	12.4	21.9	6.3	8.6
Significant	17.4	23.9	7.1	10.2
	Optimal			
Original	15.1	24.2	8.5	6.4
Significant	14.0	24.5	7.1	12.1

Discussion

The substantive results from this study can be compared to the results of similar studies that have evaluated the effect of coaching. Earlier work suggests point estimates for an overall coaching effect across both math and verbal sections of the SAT of roughly 25 points. In particular, we base our comparison on the work of Briggs (2001), Powers and Rock (1999), and Hansen (2004). The studies by Powers & Rock and Hansen used samples similar to our POP1 group, so comparisons should only be made directly to our POP1 findings. In contrast, the study by Briggs used data from NELS:88, which had the same structure and sampling design as ELS:02. So for this study coaching effect comparisons can be made for both POP1 and POP2 samples. The results from the individual studies are shown alongside our results in Table 9. In general, our estimates for the effect of coaching are similar to those of the earlier studies. Comparing effects based on regression-based estimates over time suggests that coaching has become slightly less effective for both sections of the SAT, at least on the math section of the exam. Comparing the different methodologies, Hansen's PSM analysis (which employed a "full matching" approach) produces estimates that are noticeably different for the math and verbal when compared to the others.

One interesting finding when comparing our results to those from the Briggs study using NELS:88 data is that the effect of coaching for students who have not taken the PSAT (POP2 sample) is about 10 points higher in math and 5 points higher in verbal. The difference in coaching effects in math for students in the POP1 and POP2 samples merits closer attention. On the one hand, this may indicate that coaching has a larger effect for those who have not previously taken the PSAT. On the other hand, this may be an artifact of uncontrolled confounding because we are missing information on differences in test-taking ability captured by PSAT scores. As a check on this, we re-ran the POP1 regressions after excluding the PSAT variables as controls. The resulting coaching effect in math increased from 11 to 16 points, much closer to the 22 point effect found for the POP2 sample. This leads us to believe that the higher effect estimates for the POP2 sample must be taken with a grain of salt.

In this study, the effects estimated on the basis of propensity score subclassification are more comparable to the regression results than they are with the optimal pair matching results. A key reason that coaching estimates deriving from the subclassification approach and LR approaches are so similar in this example is that the estimated propensity score distributions for coached and uncoached were not severely imbalanced at the outset. After matching through subclassification, relatively few students were excluded when estimating coaching effects. The linear regression and subclassification results use similar numbers of students, 1,552 and 1,494 for POP1 in the LR and subclassification respectively and 2,389 and 2,385 for POP2. The optimal match results are based upon far fewer students, only 706 and 894 for POP1 and POP2. This is one potential drawback to the pair matching approach¹⁰.

Properly specifying a selection function can be a challenging part of a PSM analysis. In our view, the practice of removing variables from the selection function on the basis of their statistical significance is problematic. If theory dictates that a variable be included in the selection function, lack of statistical significance is an idiosyncratic reason for removal, one that encourages "fishing for significance." It is

Table 9. Coaching Effect Estimates from Various Studies

Study	SATM		SATV	
Powers & Rock (1999)-LR	18		6	
Powers & Rock (1999)-PSM	15		6	
Hansen (2004)-PSM	23		0	
	POP1	POP2	POP1	POP2
Briggs (2001)-LR	15	8 ¹	6	1 ¹
Current Study				
Regression	11	22	6	8
Subclassification	12	22	6	9
Optimal Pair (w/ smoothing)	15	24	9	6

¹ These figures were not reported in the Briggs (2001) article but come from the same study.

also problematic because it can create a tautological case for balance. That is, balance is demonstrated only after the symptoms of imbalance—non-significant variables in the selection function—have been removed.

Going back to Hansen’s question, will matching make regression obsolete as a statistical model for causal inference? We think not. Fundamentally, it does not seem any easier to model selection as opposed to outcome: both techniques depend on the quality of the available variables to capture sources of confounding. Linear regression is also less time-consuming to implement and makes checking for interactions much easier. However, there are a number of aspects of PSM which we find appealing. The emphasis that a PSM approach places on estimating causal effects on the basis of comparable units is important. When confronted with observational data in the context of a regression analysis, a fundamental lack of comparability between treatment and control units can be easily swept under the hood. In such a context Rubin has warned that “inferences for the causal effects of treatment on such a unit cannot be drawn without making relatively heroic modeling assumptions involving extrapolations. Usually, such a unit should be explicitly excluded from the analysis” (Rubin, 2001, p. 180). Furthermore, the sensitivity analysis approach suggested by Rosenbaum (2002) offers a powerful way of analyzing the robustness of one’s results to hidden biases, and such an approach cannot be (or at least has not been) readily applied to linear regression. Conclusions based on the results of a PSM analysis can be judged relative to the robustness of the results in the presence of hidden biases. While more work is necessary to understand what does and does not constitute robustness using this method in educational research, such an approach is conceptually appealing.

Rubin (2006) has suggested that there is also an ethical advantage to the use of PSM to estimate causal effects in that all matching can be done without reference to the outcomes. In principle this would seem to support the objectivity of causal inferences. In our view this is overselling the approach. Drawing causal inferences about student achievement in an observational setting fundamentally requires us to make an informed hypothesis about (a) why subjects choose to participate (i.e., self-select) in treatment and control groups, and (b) what characteristics of these subjects are associated with the outcome of interest. All statistical modeling that follows hinges upon the quality of this foundation. If the foundation has been well-established, then just as with PSM, the estimation of an aggregate causal effect using LR happens only once. This is precisely the approach that was taken in the present study. However, if the foundation has not been well-established—and we suspect this is the case whenever covariates are being chosen on the basis of statistical significance—then there is no magic that will pull the causal rabbit out of the observational hat.

References

- Abadie, A. & Imbens, G.W. (2006). On the Failure of the Bootstrap for Matching Estimators. *National Bureau of Economic Research Technical Working Paper Series*, 325.
- Bowers, J., Fredrickson, M., & Hansen, B. (2008). RItools: Randomization Inference Tools. R package version 0.1-3
- Briggs, D. (2001). The effect of admissions test preparation: Evidence from NELS-88. *Chance*, 14, 10-18.

- Briggs, D. (2002). SAT Coaching, Bias and Causal Inference. Unpublished doctoral dissertation, University of California, Berkeley.
- Briggs, D. (2004). Evaluating SAT Coaching: Gains, Effects, and Self-Selection. In R. Zwick (Ed.), *Rethinking the SAT: The Future of Standardized Testing in University Admissions* (p. 217-234). New York: RoutledgeFarmer.
- Caliendo, M. & Kopeinig, S. (2008). Some practical guidance for the implementation of propensity score matching. *Journal of Economic Surveys*, 22 (1), 31-72.
- Callahan, R., Wilkinson, L., Muller, C., & Frisco, M. (2009). ESL Placement and Schools: Effects on Immigrant Achievement. *Educational Policy*, 23 (2), 355-384.
- College Board. (2009). *SAT Reasoning Test*. Retrieved March 21, 2009 from <http://www.collegeboard.com/student/testing/sat/about/SATI.html>.
- Frank, KA., Sykes, G., Anagnostopoulos, D., Cannata, M., Chard, L., Krause, A. (2008). Does NBPTS certification affect the number of colleagues a teacher helps with instructional matters? *Educational Evaluation And Policy Analysis*, 30, 3-30.
- Gastwirth, J., Krieger, A., & Rosenbaum, P. (2000). Asymptotic seperability in sensitivity analysis. *Journal of the Royal Statistical Society: Series B*, 62 (3), 545-555.
- Hansen, B. (2004). Full matching in an observational study of coaching for the SAT. *Journal of the American Statistical Association*, 99 (467), 609-618.
- Hansen, B. & Bowers, J. (2008). Covariate Balance in Simple, Stratified and Clustered Comparative Studies. *Statistical Science*, 23, 219-236.
- Ho, D., Imai, K., King, G., & Stuart, E. (2004). *Matchit: Matching as Nonparametric Preprocessing for Parametric Causal Inference*. Available from <http://gking.harvard.edu/matchit>
- Ho, D., Imai, K., King, G., & Stuart, E. (2007). Matching as nonparametric preprocessing for reducing model dependence in parametric causal inference. *Political Analysis*, 15 (3), 199-236.
- Holland, P. (1986). Statistics and causal inference. *Journal of the American Statistical Association*, 81, 945-960.
- Hong, G. & Raudenbush, S. (2005). Effects of kindergarten retention policy on children's cognitive growth in reading and mathematics. *Educational Evaluation and Policy Analysis*, 27 (3), 205-224.
- Hong, G., & Raudenbush, S. (2006). Evaluating kindergarten retention policy: A case study of causal inference for multilevel observational data. *Journal of the American Statistical Association*, 101 (475), 901-910.
- Keele, L. (2008). rbounds: Perform rosenbaum bounds sensitivity tests for matched data. R package version 0.1
- Lawrence, I., Rigol, G., Essen, TV., & Jackson, C. (2004). A historical perspective on the content of the SAT. In R. Zwick (Ed.), *Rethinking the SAT: The Future of Standardized Testing in University Admissions* (p. 57-74). New York: RoutledgeFarmer.
- Morgan, S. (2001). Counterfactuals, causal effect heterogeneity, and the catholic school effect on learning. *Sociology of Education*, 744, 341-374.
- Powers, D. & Rock, D. (1999). Effects of coaching on SAT I: Reasoning test scores. *Journal of Educational Measurement*, 36 (2), 93-118.
- R Development Core Team. (2008). R: A language and environment for statistical computing. Available from <http://www.R-project.org>
- Rosenbaum, P. (2002). *Observational studies* (2nd ed.). New York: Springer.
- Rosenbaum, P. & Rubin, D. (1983). The central role of the propensity score in observational studies for causal effects. *Biometrika*, 70 (1), 41-55.
- Rubin, D. (1976). Inference and missing data. *Biometrika*, 63, 581-592.
- Rubin, D. (2001). Using propensity scores to help design observational studies: Applications to the tobacco litigation. *Health Services & Outcomes Research Methodology*, 2, 169-188.
- Rubin, D. (2006). The design versus the analysis of observational studies for causal effects: Parallels with the design of randomized trials. *Statistics in Medicine*, 26 (1), 20-36.
- Ruhm, K. & Waldfogel, J. (2007). Does prekindergarten improve school preparation and performance? *Economics of Education Review*, 26, 33-51.

Appendix: Variable Index

1. *SATM*: Score on the SATM. Based on ELS:02 [txsatm].
2. *SATV*: Score on the SATV. Based on ELS:02 [txsatv].
3. *PSATM*: Score on the PSATM. Based on ELS:02 [f1rpsatm].
4. *PSATV*: Score on the PSATV. ELS:02 [f1rpsatv].
5. *AGE*: Approximate age in months at time of first follow-up survey. Modification of ELS:02 [bydob_p].
6. *SES*: The SES index created for ELS. Based on ELS:02 [byses2].
7. *FEMALE*: Dummy variable indicating respondent is female. Modification of ELS:02 [bysex].
8. *RACE*: Composite variable indicating whether respondent is black or a Native American (including natives of Alaska and Hawaii). Based on ELS:02 variables byrace_2, byrace_4, and byrace_5.
9. *ASIAN*: Dummy variable indicating respondent is Asian. Based on ELS:02 [byrace_3].
10. *HISPANIC*: Dummy variable indicating respondent is Hispanic. Based on ELS:02 [bys15].
11. *PRIVATE*: Dummy variable indicating whether respondent attends private school. Modification of ELS:02 [bysctrl].
12. *RURAL*: Dummy variable indicating whether respondent attended a rural school. Modification of ELS:02 [byurban].
13. *URBAN*: Dummy variable indicating whether respondent attended an urban school. Modification of ELS:02 [byurban].
14. *AP*: Dummy variable indicating whether respondent took an AP or IB class. Modification of ELS:02 [f1rapib].
15. *REM_ENG*: Dummy variable indicating whether respondent took a remedial English class. Based on ELS:02 [bys33d].
16. *REM_MATH*: Dummy variable indicating whether respondent took a remedial Math class. Based on ELS:02 [bys33e].
17. *COLL_PREP*: Dummy variable indicating whether respondent's high school program was college preparatory. Modification of ELS:02 [byschprg].
18. *BYREAD*: Student's score on the ELS:02 administered Base Year reading test. Based on ELS:02 [bytxrstd].
19. *BYMATH*: Student's score on the ELS:02 administered Base Year math test. Based on ELS:02 [bytxmstd].
20. *MATH_CRD*: Number of math credits from the student's transcript around the time of the first follow-up survey. Based on ELS:02 [f1rhma_c].
21. *HW*: Dummy variable indicating whether respondent spent more than 10 hours/week on homework. Modification of ELS:02 [bys34b].
22. *ESL*: Dummy variable indicating that respondent speaks English as a second language. Modification of ELS:02 [bys67].
23. *EDU_AFTER_HS*: Dummy variable indicating whether respondent planned to continue education immediately after high school. Modification of ELS:02 [bys57].
24. *COLLEGE_INFO*: Dummy variable indicating that a respondent did not seek college information from any of the listed sources (parents, counselors, etc). Based on ELS:02 [bys59k].
25. *PRNTS_DISC_PREP*: Dummy variable indicating whether respondent discussed SAT/ACT preparation with parents often. Modification of ELS:02 [bys86f].
26. *PRNTS_DISC_SCH*: Dummy variable indicating whether respondent discussed school courses with parents often. Modification of ELS:02 [bys86a].
27. *COACH*: Dummy variable indicating whether respondent received coaching for the SAT/ACT. Modification of ELS:02 [f1s22b].
28. *UNDERPERFORM.M*: Dummy variable indicating that a student may have underperformed on the Base Year Math test relative to their GPA. Defined further in the data section.
29. *UNDERPERFORM.V*: Dummy variable indicating that a student may have underperformed on the Base Year Reading test relative to their GPA. Defined further in the data section.
30. *NERVES.M*: Dummy variable indicating that a student's PSATM score was much lower than anticipated based on the Base Year Math score, possibly due to nervousness. Defined further in the data section.
31. *NERVES.V*: Dummy variable indicating that a student's PSATV score was much lower than anticipated based on the Base Year Reading score, possibly due to nervousness. Defined further in the data section.
32. *HI_SES*: Dummy variable indicating that a student's value on the ELS ses variable was in the top quarter for all students in the survey. Modification of ELS:02 [byses2].

Notes

1. In particular, his methodology reduced pre-matching imbalances between treatment and control subjects substantially in 27 covariates and still allowed more of the data to be used than would have been possible in a regression analysis (Hansen, 2004, p. 617).
2. The PSAT is essentially a pre-test for the SAT taken by grade 11.
3. “Substantially worse” here is defined as having a PSAT score less than two RMSEs below their predicted score based on the ELS:02 Base Year test.
4. The variables with the bulk of the missing data were ASIAN, BLACK, NATIVE, HISPANIC, REM_ENG, REM_MATH, COLLEGE_INFO, PRNTS_DISC_PREP, and PRNTS_DISC_SCH
5. For those variables with this extra level, the variable name shown in later tables was modified to include a postscript of “1” or “NA”. For example, “REM_ENG” becomes either “REM_ENG1” or “REM_ENGNA”.
6. It is worth noting that this choice often appears rather fluid in empirical applications of PSM. Because there are few guiding principles that inform the number of subclasses, some researchers essentially use this as a variable that can be manipulated to demonstrate that sufficient balance has been obtained.
7. We have taken a parametric approach to estimate standard errors for the purpose of conducting tests of significance. These standard errors are presented in Table 4. One potential alternative to this parametric approach would be the bootstrap, but this was shown to be a poor choice for matched data (Abadie and Imbens, 2006). Ho et al. (2007) take the position that since pretreatment variables are typically assumed to be fixed and exogenous, standard parametric adjustments are appropriate. Further discussion of this issue is outside the scope of the present study.
8. For example, see , as in Hong & Raudenbush (2006). Though it is not made explicit in the paper that a stepwise selection method was used, this was indicated to us in a personal correspondence (G. Hong, Personal Communication, September 8, 2008).
9. In the terminology of Rosenbaum (2002), these ranges correspond to a Γ of 1.3.
10. An alternative to pair matching that still preserves the notion of finding specific matches for each treatment unit would be fixed ratio matching or variable matching.

Send correspondence to: Ben Domingue
 University of Colorado
 Email: benjamin.domingue@colorado.edu

This research was supported by a grant from the American Educational Research Association which receives funds for its “AERA Grants Program” from the National Sciences Foundation and the National Center for Education Statistics of the Institute of Education Sciences (U.S. Department of Education) under NSF Grant #DRL-0634035. Opinions reflect those of the authors and do not necessarily reflect those of the granting agencies.

Hierarchically Ordered Instructional Activities for the Development of Clinically Relevant Competencies: A Sequential Path Model

Frank J. Papa

University of North Texas Health Science Center

Randall E. Schumacker

University of Alabama

A potential benefit in medical education is the use of hierarchially ordered instruction (Bloom's taxonomy of educational objectives) to progressively develop the intellectual capabilities leading to clinically relevant competencies in medical students. In this study, faculty observations indicated significant improvements in the students clinically relevant competencies compared to students trained without hierarchially ordered instructional methods. Student responses demonstrated that the sequential ordering of instructional activities significantly contributed to the development of their evolving clinical competencies ($R^2_M = .62$, $p < .05$).

Throughout the 20th century, developmental psychologists observed that human intellectual capabilities and knowledge-based competencies matured in a roughly developmentally progressive manner. That is, in a series of hierarchically ordered stages representing a logical progression from simple to more complex capabilities and competencies. Bloom (1956) codified these observations by defining six hierarchically ordered steps beginning with information/knowledge acquisition, followed by comprehension, then application, analysis, synthesis, and finally evaluation oriented capabilities and competencies. He further posited that hierarchically ordered instruction (i.e., instruction designed to sequentially develop these six capabilities/competencies – one step at a time) represented an effective instructional strategy (Bloom, 1956, p. 78-88).

While there is reason to believe that medical students could benefit from hierarchially ordered instructional activities, there is little if any evidence of its utility in the medical education setting. Our research therefore investigated the use of hierarchically ordered instruction as a means of improving the clinically relevant competencies of sophomore medical students in the area of respiratory medicine. The clinically relevant competencies of particular interest in this study involved two core knowledge-based competencies: diagnostic ability and ability to use basic and clinical science concepts and principles to explain clinical phenonema. Our approach involved faculty observations of the performance of students trained with hierarchically ordered instructional activities compared with the performance of students trained under our traditional curriculum. Our approach requires a brief review of the traditional instruction and assessment methods in use throughout our year two curriculum at the time this research was conducted and a review of the specific instruction and assessment methods used in our sequential ordered instructional program.

Traditional Instruction and Assessment Methods

No formalized learning sciences principles, methods, or framework served as a guideline or blueprint for the construction of our traditional Year Two coursework. A review of scheduled curricular activities revealed that approximately 85 - 90% of all faculty/student interactions occurred in the form of traditional lecture based instruction, that is, information acquisition-oriented instruction in a large classroom environment. Student and faculty surveys demonstrated that approximately 10% - 15% of scheduled class time involved the use of case vignettes or approximately 1 – 2 case vignettes per contact hour. While case vignettes represented a small component of instructional activities, there did not appear to be any evidence that they were formally used to develop higher level intellectual capabilities (such as comprehension and application capabilities). There was no evidence that a formal hierarchical ordering of instruction was in place, that is, passive information acquisition oriented instruction preceding case vignette based comprehension activities, followed in turn by case vignette based application activities.

Following a review of our Year Two examinations, approximately 85% of the test items appeared to be first-order, multiple choice, recognition oriented items noted by Bloom (1956) as tests of the recall of acquired information. Approximately 10% - 15% of test items were in the form of more clinically relevant second-order questions (case vignette based). A second-order question was defined as students being presented with a case vignette followed by a basic or clinical science oriented question. Students used the signs and symptoms in the vignette to draw an inference regarding the disease represented by the

vignette. Having drawn a diagnostic inference, the student then used their diagnoses as the basis for answering the basic or clinical science oriented question associated with the test item vignette. Thus, a second-order case vignette test item provided students an opportunity to demonstrate both their diagnostic competency and evolving ability to apply various basic and clinical science concepts most relevant to the disease represented by the vignette.

Sequential Ordered Instruction Method

A sequential, hierarchically ordered instructional method was developed for use in a respiratory medicine course in year two (RM2). The new approach replaced all traditional passive in-class instructional methods with active, hierarchically ordered instructional activities. The RM2 course contained eight problem specific instructional modules as the basis for reorganizing the content of the course. The eight problems covered were adult dyspnea, pediatric dyspnea, hemoptysis, cough, epistaxis, rhinitis, otalgia and sore throat. Each problem specific instructional module began with information acquisition (pre-classroom reading assignments oriented to the problem at hand and its disease differentials), followed by in-class, comprehension oriented instructional activities, followed by in-classroom, application instructional activities. The RM2 course used reality oriented case-vignettes as the basis for all comprehension and application oriented instructional activities and used audience response technologies to support student engagement and teacher/student interaction in all classroom based comprehension and application oriented instructional activities. The RM2 course also used a web-based artificial intelligence tool (called KBIT) to provide problem-specific, hierarchically ordered differential diagnostic instruction to students. The number of case vignettes used in both the classroom and via computer-based training was sufficient to develop and assess the attainment of a rudimentary level of competence in diagnosing the disease differentials for each of the eight core respiratory problems as well as explain the basic and clinical science concepts underlying each problem and its disease etiologies.

RM2 scheduled coursework was allocated as follows. Thirty percent (30%) was devoted to information acquisition in the form of student self study of faculty defined reading assignments and learning objectives to guide them in their self study efforts. Sixty percent (60%) was devoted to a blend of comprehension and application-oriented instructional activities for each of the eight problem-specific modules. In developing student comprehension of key concepts and principles, faculty were asked to use the 5–6 case vignettes (per hour) to model how they themselves might use their diagnostic and explanatory knowledge in each case. Faculty were also encouraged to engage students in rudimentary attempts to apply their evolving intellectual processes and knowledge base to the diagnosis and explanation of case-related issues. Ten percent (10%) was devoted purely to Problem-Specific Knowledge Base Application Sessions (PSKBAS). These PSKBAS activities were held at the end of the course.

Method and Procedures

Participants

In-class observations by nine senior faculty ($n = 9$) served as the basis for assessing the students' evolving diagnostic, explanatory, and overall clinical competencies. These observations occurred in the eight PSKBAS which were held at the end of RM2 and conducted in a large classroom setting. A total of eight PSKBAS sessions occurred with each session covering one of the eight previously noted respiratory problems. The format for each PSKBAS involved the use of approximately 20 case-based vignettes per hour as stems for questions designed to elicit the students' ability to correctly diagnose each case vignette and explain the pathologic, pathophysiologic, physiologic, biochemical, or pharmacologic issues associated with particular aspects of each case vignette (second order test items). Up to four faculty members, representing various disciplines such as internal medicine, pediatrics, surgery, pathology, biochemistry, physiology, or pharmacology participated in each of the eight PSKBAS.

PSKBAS focusing upon adult dyspnea, pediatric dyspnea, hemoptysis and cough involved two hours of second-order case-vignette based questions. The remaining four problems (epistaxis, rhinitis, otalgia and sore throat) involved one hour of these second-order case vignette presentations. During these 12 hours of PSKBAS, over 200 hundred case vignettes with clinically related questions were presented to the students. Faculty observed the answers students gave for each case vignette via audience response technologies that collected and projected in real time the summarized reports of their responses on a large projection screen. Shortly following the completion of the RM2 course, a ten item questionnaire was

presented to the nine senior faculty members primarily responsible for constructing and directing all of the RM2 course comprehension and application sessions as well as the eight PSKBAS.

The faculty observations served as a means of directly evaluating the evolving clinical capabilities and competencies of the RM2 trained students, and indirectly, the effectiveness of learning sciences framework used to design RM2 coursework. Three faculty questionnaire items were specifically designed to gather faculty impressions of student performance in regards to their overall clinical capabilities, diagnostic capabilities, and explanatory capabilities (ability to use basic and clinical science concepts and principles to explain various clinical phenomena). These three questions all had the same answer set which enabled faculty to compare the clinical capabilities of students in the RM2 course to students in our traditional instructional course. A chi-square analysis was applied to these faculty based observations and impressions.

One hundred and nine students ($n = 109$) out of 121 students completed a 55 item post RM2 course questionnaire. Students responded to Likert scaled items consisting of five possible responses (strongly agree, agree, neutral, disagree, strongly disagree). The first seventeen items in Table 1 were placed in sequential ordered sets. The eighteenth item was used as the dependent variable. The item sets and dependent variable were analyzed in a sequential path model.

Research Questions

Our research investigation into the development of clinically relevant competencies resulted in two research questions. Will clinical faculty observations demonstrate that the performance of RM2 trained students exceeded students trained previously via our traditional curriculum? This research question involved three areas of inquiry:

- a) Overall level of clinical competence achieved
- b) Diagnostic capabilities
- c) Level of explanatory capabilities achieved (the ability to use basic and clinical sciences concepts and principles to explain clinical phenonema).

Our second research question attempted to determine if the development of clinically relevant diagnostic competencies begins with information acquisition, followed by comprehension based capabilities, followed by an ability to successfully apply their evolving knowledge base. Will student responses provide evidence of a sequential ordered instructional change in clinical competencies?

Results

For the nine (9) clinical faculty responses, a 3 X 2 chi-square analysis was performed for all three questions posed to faculty. For *level of clinical competence achieved*, 78% of faculty reported that the RM2 trained students overall level of clinical competence exceeded expectations, which was statistically significant ($\chi^2 = 8.686$, $df = 2$, $p = 0.013$). For *diagnostic capabilities*, 67% of faculty reported that the RM2 students diagnostic capabilities exceeded expectations, which was statistically significant ($\chi^2 = 5.992$, $df = 2$, $p = 0.050$). For *level of explanatory capabilities achieved*, 75% of responding faculty reported that the ability of RM2 students to explain phenomena via use of constructs from basic and clinical sciences exceeded expectations, which was statistically significant ($\chi^2 = 7.014$, $df = 2$, $p = 0.030$). An effect size was computed to measure the degree of departure from the null hypothesis in standard units. The effect size (r) = 0.257 and Cohen's $d = 0.532$ (unbiased), which indicated a medium effect size (Schumacker & Akers, 2001).

For the one hundred nine student responses ($n = 109$) to the questionnaire items in Table 1, there were 16 missing data points out of 2,180 (109 subjects x 18 variables) representing .007 percent missing data. Missing values were replaced with a 0 (zero) so that the summative scores for the independent variables would not be influenced by mean substitution or replacement by a raw score scale value of 3 (neutral) giving more importance to a neutral item response.

Three sequential ordered independent variables were then created by summing the respective questionnaire items. Student questionnaire items 1, 2, 3, 4, 5, 6, 7, 8 were summed to represent information acquisition oriented instructional activities (Acquisition); items 9, 10, 11, 12 were summed to represent comprehension oriented instructional activities (Comprehension); and items 13, 14, 15, 16, 17

Table 1. Sequential Ordered Sets of Questionnaire Items

Course Component	Questionnaire Items
Information Acquisition	<p>#1: The printed objectives for Pathology's respiratory and ENT reading assignments were sufficiently detailed to help me focus on the most important concepts.</p> <p>#2: The printed objectives set forward for Pathology's respiratory and ENT chapters enhanced my performance on examinations covering these chapters.</p> <p>#3: I understood most of the concepts contained within Nelson's textbook on my own with little need for clarification from an instructor.</p> <p>#4: The printed objectives for Pediatric's respiratory and ENT reading assignments were sufficiently detailed to help me focus on the most important concepts.</p> <p>#5: The printed objectives set forward for Pediatric's chapters enhanced my performance on examinations covering these chapters.</p> <p>#6: I understood most of the concepts contained within Internal Medicine's textbook on my own with little need for classroom sessions.</p> <p>#7: The printed objectives for Internal Medicine's respiratory reading assignments were sufficiently detailed to help me focus on the most important concepts.</p> <p>#8: The printed objectives set forward for Internal Medicine's reading assignments enhanced my performance on examinations covering these chapters.</p>
Comprehension	<p>#9: Interactive class sessions enhanced my understanding of the self study materials.</p> <p>#10: Interactive sessions following, and keyed to prior self study of a required reading assignments) should be adopted by other System Two course faculty.</p> <p>#11: KBIT's WEB-based 'training' cases (cases that did not require a diagnosis) helped me better understand the variety of ways patients with a given disease could present in the clinical arena.</p> <p>#12: The type of knowledge placed within an 'explanatory' knowledge base is related to but organized in a different format from 'diagnostic' knowledge.</p>
Application	<p>#13: Following a wrong diagnosis, KBIT's case-specific 'feedback' enhanced my understanding of how signs and symptoms could have been better used to correctly diagnose the case.</p> <p>#14: Spending an hour or more focusing on a specific clinical problem (during application sessions) improved my understanding of how to initiate my approach to cases presenting with that particular clinical problem.</p> <p>#15: Compared with traditional lectures, the application sessions provided a greater opportunity for me to integrate textbook and lecture derived knowledge about clinical problems and their disease differentials</p> <p>#16: The application sessions enhanced my evolving patient care management capabilities.</p> <p>#17: Application session cases and questions engaged me and challenged me in ways different than traditional course examinations.</p>
Dependent Variable	<p>#18: My overall level of clinical competence (i.e., diagnostic capabilities, clinical skills, treatment skills, patient care management skills) improved in Respiratory Medicine Two (RM2).</p>

were summed to represent application oriented instructional activities (Application). These three sequential ordered variables were used to test the sequential relationship developed for the three core instructional methods used in RM2 (Acquisition, Comprehension, Application) as related to the dependent outcome variable (improved clinical competencies, item 18).

Table 2 indicates the variable scale statistics including Cronbach internal consistency reliability coefficients for the three independent variables.

Table 3 indicates the correlation amongst the variables in the path model. The correlation between Acquisition and Application was the only non-significant correlation coefficient. In a manner consistent with Bloom's taxonomy, the students' sense of their improved clinical competencies (Y) indicated increased correlation values across the sequential ordered variables ($r = 0.176$ to 0.388 to 0.547).

Table 2. Variable Descriptive Statistics (*N* = 109)


Variable	Items	Mean	SD	Reliability
Acquisition	8	21.06	5.14	.786
Comprehension	4	8.62	2.94	.606
Application	5	12.61	4.770	.839
Y	1	2.64	1.03	n/a

Table 3. Correlation Matrix (*N* = 109)

Variable	Y	Acquisition	Comprehension	Application
Y	1.000			
Acquisition	0.176*	1.000		
Comprehension	0.388**	0.430**	1.000	
Application	0.547**	0.075	0.568**	1.000

* $p < 0.05$, $df = 108$ (one-tail test); ** $p < 0.01$, $df = 108$ (one-tail test)

A path model analysis using structural equation modeling software (LISREL) statistically tested the research question: Will student responses provide evidence of improvements in clinical competencies due to sequentially ordered instruction? The path model was diagrammed as:



The correlation matrix is in Table 3. The sequential ordered variables yielded statistically significant path coefficients and R-squared values. The path model chi-square value was non-significant at the $p < .01$ level of significance indicating that the data fit the hypothesized theoretical path model ($\chi^2 = 8.79$, $df = 3$, $p = 0.03$). The overall R^2 value for this path model, which included direct and indirect effects, was $R^2_m = 0.615$. Findings therefore indicated that 62 percent of the variation in Y (improved clinical competencies) was explained by the sequential nature of the independent variables direct and indirect effects.

Conclusions

The curricular reform initiative represented in this study was research designed to improve the clinical capabilities and competencies of medical students. The authors believed that a learning sciences framework utilizing hierarchically ordered instructional sequencing (information acquisition, followed by comprehension, followed by application activities) could improve the clinically relevant competencies of sophomore medical students. Faculty observations provided evidence of performance improvements in the students' overall clinical capabilities and more specifically in their diagnostic and explanatory capabilities. Student feedback also validated the underlying learning sciences framework used in the design of the RM2 course and provided a cognitively based explanation as to why the students overall clinical competencies improved.

The path analysis results supported the research question, that is, intellectual development (intellectual processes and knowledge based capabilities) evolved in a logical progression as indicated by a statistically significant sequential ordering of the independent variables: Acquisition, Comprehension, and Application. The path analysis results therefore suggest that hierarchically ordered instruction represents useful means of supporting the evolving nature of the intellectual processes and knowledge based capabilities of medical students.

Efforts intended to champion curricular reform initiatives require a well planned, multi-faceted approach to evaluating the outcomes of the proposed instructional interventions (Hardin, Grant, Buckley, & Hart, 1999; Wilkes & Bligh, 1999; Norman, Eva, & Schmidt, 2005; Mamede, Schmidt, & Norman, 2006). Without such an approach, it is unlikely that curricular innovators will be able to provide evidence sufficient to argue that their particular instructional interventions merit further consideration, refinement and perhaps emulation. The results of this study provided evidence that a larger scale implementation of our sequential ordered instruction could yield positive performance improvements throughout coursework

representing the entire second year of the curriculum. A recent investigation involving the use of hierarchically ordered instructional sequencing (and other learning sciences principles) across our year two curriculum (called an Application-Oriented Curriculum) has provided evidence of significant improvement of students on national Licensing Board examinations (Papa et al., in review).

References

- Bloom B. S. (1956). *Taxonomy of Educational Objectives: The classification of educational goals. Handbook I: Cognitive Domain*. New York, NY. David McKay Company, Inc.
- Norman G R., Eva K.W., & Schmidt H.G. (2005). Implications of psychology-type theories for full curricular interventions. *Medical Education*, 39, 247-249.
- Mamede, S., Schmidt H.G., & Norman G.R. (2006). Innovations in Problem-based Learning: What can we learn from recent studies? *Advances in Health Sciences Education*, 11, 403-422.
- Hardin R.M., Grant J., Buckley G., & Hart I.R. (1999). BEME Guide No 1: Best Evidence Medical Education. *Medical Teacher*, 21 (6), 553-662.
- Papa, F.J., Oglesby, M.W., Dubin, B.D., Gwartz, P.A., Martin, M.W. & Putthoff, S.L. (in review). An Application-Oriented Curriculum: Rationale, Learning Sciences Framework and Outcomes. *Medical Education*.
- Wilkes M. & Bligh J. (1999). Evaluating educational interventions. *British Medical Journal*, 318,1269-1272.

Send correspondence to: Frank J. Papa
University of North Texas Health Science Center
Email: fpapa@hsc.unt.edu

Multiple Linear Regression Viewpoints ***Information for Contributors***

Multiple Linear Regression Viewpoints (MLRV) is a publication sponsored by the American Educational Research Association's Special Interest Group on Multiple Linear Regression: The General Linear Model (MLR/GLM SIG). It is published twice a year to facilitate communication among professionals who focus their research on the theory, application, or teaching of multiple linear regression models and/or the general linear model. Manuscripts submitted to *MLRV* should conform to the language, style, and format of the Publication Manual of the American Psychological Association (5th ed., 2001). Three copies (two blind) of a doubled spaced manuscript (including equations, footnotes, quotes, and references) of approximately 25 pages in length, a 100 word abstract, and an IBM formatted diskette with the manuscript formatted in WordPerfect or Word should be submitted to one of the editors listed below.

Mathematical and Greek symbols should be clear and concise. All figures and diagrams must be photocopy-ready for publication. Manuscripts will be anonymously peer reviewed by two editorial board members. Author identifying information should appear on the title page of only one submitted manuscript. The review process will take approximately 3 months. A letter acknowledging receipt of the manuscript will be sent to the first author, and upon review, a letter indicating the peer review decision will be sent to the first author. Potential authors are encouraged to contact the editors to discuss ideas for contributions or determine if their manuscript is suitable for publication in *MLRV*.

EDITOR INFORMATION

Randall E. Schumacker, Editor *MLRV*
College of Education
P.O. Box 870231
Carmichael Hall 316
The University of Alabama
Tuscaloosa, AL 35487-0231
(205) 348-6062
rschumacker@ua.edu

T. Mark Beasley, Associate Editor
Department of Biostatistics
School of Public Health
309E Ryals Public Health Bldg.
University of Alabama, Birmingham
Birmingham, AL 35294
(205) 975-4957
mbeasley@uab.edu

ORDER INFORMATION

Cynthia Campbell, MLR/GLM SIG Executive Secretary
Department of Educational Technology, Research & Assessment
Northern Illinois University
DeKalb, IL 60115-2854
ccampbell@niu.edu

Check out our website at: <http://mlrv.ua.edu/>

POSTMASTER: Send address changes to:
Cynthia Campbell, MLR/GLM SIG Executive Secretary
Department of Educational Technology, Research & Assessment
Northern Illinois University
DeKalb, IL 60115-2854

Multiple Linear Regression Viewpoints (ISSN 0195-7171) is published by the
AERA Special Interest Group on Multiple Linear Regression: General Linear Model
through the **University of Alabama** and the **Northern Illinois University**.