



Regression with imputed covariates: A generalized missing-indicator approach

Valentino Dardanoni, Salvatore Modica, Franco Peracchi

► To cite this version:

Valentino Dardanoni, Salvatore Modica, Franco Peracchi. Regression with imputed covariates: A generalized missing-indicator approach. *Econometrics*, 2011, 10.1016/j.jeconom.2011.02.005 . hal-00815561

HAL Id: hal-00815561

<https://hal.science/hal-00815561>

Submitted on 19 Apr 2013

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Accepted Manuscript

Regression with imputed covariates: A generalized missing-indicator approach

Valentino Dardanoni, Salvatore Modica, Franco Peracchi

PII: S0304-4076(11)00032-7

DOI: [10.1016/j.jeconom.2011.02.005](https://doi.org/10.1016/j.jeconom.2011.02.005)

Reference: ECONOM 3451

To appear in: *Journal of Econometrics*

Received date: 2 October 2009

Revised date: 11 December 2010

Accepted date: 8 February 2011



Please cite this article as: Dardanoni, V., Modica, S., Peracchi, F., Regression with imputed covariates: A generalized missing-indicator approach. *Journal of Econometrics* (2011), doi:10.1016/j.jeconom.2011.02.005

This is a PDF file of an unedited manuscript that has been accepted for publication. As a service to our customers we are providing this early version of the manuscript. The manuscript will undergo copyediting, typesetting, and review of the resulting proof before it is published in its final form. Please note that during the production process errors may be discovered which could affect the content, and all legal disclaimers that apply to the journal pertain.

REGRESSION WITH IMPUTED COVARIATES: A GENERALIZED MISSING-INDICATOR APPROACH*

VALENTINO DARDANONI, SALVATORE MODICA, AND FRANCO PERACCHI

ABSTRACT. A common problem in applied regression analysis is that covariate values may be missing for some observations but imputed values may be available. This situation generates a trade-off between bias and precision: the complete cases are often disarmingly few, but replacing the missing observations with the imputed values to gain precision may lead to bias. In this paper we formalize this trade-off by showing that one can augment the regression model with a set of auxiliary variables so as to obtain, under weak assumptions about the imputations, the same unbiased estimator of the parameters of interest as complete-case analysis. Given this augmented model, the bias-precision trade-off may then be tackled by either model reduction procedures or model averaging methods. We illustrate our approach by considering the problem of estimating the relation between income and the body mass index (BMI) using survey data affected by item non-response, where the missing values on the main covariates are filled in by imputations.

JEL Classification Numbers: C12, C13, C19.

Keywords Missing covariates; Imputations; Bias-precision trade-off; Model reduction; Model averaging; BMI and income.

INTRODUCTION

A common problem in applied regression analysis is that covariate values may be missing for some observations but imputed values may be available, either values provided by the data-producing agency or directly constructed by the researcher. This problem has received little attention compared to the more general problem of missing covariate values, but is of considerable practical relevance as all empirical researchers know well. In many cases, it is safe to assume that the mechanism leading to missing covariate values does not depend on the outcome of interest. In these cases, one can ignore the missing data mechanism and focus on the problem of what use to make of the available imputations.

There are two main approaches to this problem. One is to simply ignore the imputations and only use the observations with complete data on all covariates—the so-called complete-case analysis. Although this may entail a loss of precision, it has the strong appeal of yielding an unbiased estimator of the parameters of interest when the missing data mechanism is ignorable. The other approach is more concerned with precision

Date: February 14, 2011.

* Corresponding author: franco.peracchi@uniroma2.it. Affiliations: Valentino Dardanoni and Salvatore Modica, University of Palermo; Franco Peracchi, Tor Vergata University and EIEF. We thank Jan Magnus for comments and insightful discussions, and an Associate Editor and a referee for helpful suggestions.

and replaces the missing covariate values with the imputations. A refined version of this approach corrects for incorporating the imputed values by some variant of the so-called missing-indicator method (Little [12], Horton and Kleinman [7], and Little and Rubin [13]), which consists of augmenting the regression model with a set of binary indicators for each covariate with missing values. Although frequently used in practice, this approach is known to produce biased estimates (Jones [8], Horton and Kleinman [7]). It also raises the problem of how to assess precision of the estimators, a problem that we ignore in this paper because it can easily be handled by multiple imputation methods (Rubin [21]).

Thus, when covariate values are missing we face a trade-off between bias and precision: the complete cases are often disarmingly few, but replacing the missing observations with the imputed values to gain precision may lead to bias. In this paper we formalize the bias-precision trade-off by showing that one can augment the regression model with a set of auxiliary variables so as to obtain, under weak assumptions about the imputations, the same unbiased estimator of the parameters of interest as complete-case analysis. Given this augmented model, the bias-precision trade-off may then be tackled either by standard model reduction procedures or, more aptly in our view, by model averaging methods.

We illustrate our approach by considering the problem of estimating the relationship between income and the body mass index (BMI) using survey data affected by item non-response, where the missing values on the main covariates are filled in by imputation.

The sequel of the paper is organized as follows. Section 1 presents the basic notation. Section 2 discusses complete-case analysis. Sections 3 and 4 present the augmented model with auxiliary variables and discuss its missing-indicator interpretation. Section 5 contains our main result. Section 6 discusses the trade-off between bias and precision. Section 7 presents our application to modeling the relation between BMI and income. Finally, Section 8 offers some concluding remarks.

1. NOTATION

Observations are indexed by $n = 1, \dots, N$, and covariates by $k = 0, 1, \dots, K - 1$, with $k = 0$ corresponding to the constant term and $K > 1$. We consider the classical linear model

$$\mathbf{y} = \mathbf{X}\boldsymbol{\beta} + \mathbf{u}, \quad (1)$$

where \mathbf{y} is the $N \times 1$ vector of observations on the outcome of interest, \mathbf{X} is an $N \times K$ matrix of observations on the covariates, $\boldsymbol{\beta}$ is the $K \times 1$ vector of coefficients and \mathbf{u} is an $N \times 1$ vector of homoskedastic and serially uncorrelated regression errors with zero mean conditional on \mathbf{X} .

A subsample with incomplete data is a group of observations where one or more covariates are missing. Because the constant term is always observed, the number of possible subsamples with incomplete data is equal to $2^{K-1} - 1$. Not all such subsamples need be present in a data set. In addition to the subsample with complete data (indexed by $j = 0$), we assume to have $J \leq 2^{K-1} - 1$ subsamples with incomplete data, indexed by $j = 1, \dots, J$. This formulation covers both the case when some patterns of missing covariates are not present in the data and the case when the investigator decides to drop from the analysis some groups with incomplete data.

Let N_j , K_j and $K_j^* = K - K_j$ respectively denote the sample size, the number of available covariates (the covariates with no missing values, including the constant term), and the number of missing covariates in the j -th subsample. By construction $\sum_{j=0}^J N_j = N$, $K_0 = K$, $K_0^* = 0$ and $1 \leq K_j, K_j^* < K$ for $j = 1, \dots, J$. Let \mathbf{y}^j , \mathbf{X}_a^j and \mathbf{X}_m^j respectively denote the $N_j \times 1$ outcome vector, the $N_j \times K_j$ submatrix containing the values of the available covariates, and the $N_j \times K_j^*$ submatrix containing the values of the missing covariates for the j -th subsample. Also let $\mathbf{X}^j = [\mathbf{X}_a^j, \mathbf{X}_m^j]$, an $N_j \times K$ matrix. We assume that $\mathbf{X}^0 = \mathbf{X}_a^0$ is of full column rank, which implies that $N_0 \geq K$.

2. COMPLETE-CASE ANALYSIS

Our benchmark in dealing with missing values is the so called complete-case method, which uses only the observations with complete data on all covariates.

Let \mathbf{M} denote the $N \times K$ missing-data indicator matrix, whose (n, k) -th element m_{nk} takes value 1 if the n -th observation contains a missing value on the k -th covariate and value 0 otherwise. The following assumption is common to most approaches to the problem of missing covariate values and is maintained throughout this paper.

Assumption 1 (Ignorability). *\mathbf{M} and \mathbf{y} are conditionally independent given \mathbf{X} .*

By symmetry of conditional independence, it is easily seen that Assumption 1 is equivalent to the following two assumptions

$$P(\mathbf{y} \mid \mathbf{X}, \mathbf{M}) = P(\mathbf{y} \mid \mathbf{X}) \quad (2)$$

and

$$P(\mathbf{M} \mid \mathbf{y}, \mathbf{X}) = P(\mathbf{M} \mid \mathbf{X}). \quad (3)$$

Assumption 2 basically says that if we knew the true values of the missing covariates, knowing the pattern of missing data would not help in predicting \mathbf{y} . Assumption 3 implies that the missing data mechanism, seen as a function of \mathbf{y} and \mathbf{X} , depends on \mathbf{X} only. Assumption 1 may fail if, for example, observations with missing covariate values have a different regression function than observations with no missing values. On the other hand, it does not place restrictions on how \mathbf{M} is generated from \mathbf{X} . For example, \mathbf{M} may exhibit patterns such that cases with low or high levels of some covariates systematically have a greater percentage of missing values.

Theorem 1 below provides a formal proof of the fact that, under Assumption 1, the OLS estimator for the complete case is unbiased. This result has been known for long time, but may be considered a “folk theorem”. Little [12] and Little and Rubin [13] attribute it to an unpublished 1986 technical report by William Glynn and Nan Laird. Private communication with Nan Laird however informs us that the report has never been published and is no longer available. Jones [8] offers a proof for the case of two covariates, one of which has missing values, whereas Wooldridge [24], p. 553, shows that the two-stage least squares estimator for the complete case is consistent.

Theorem 1 (Complete-Case Estimation). *If Assumption 1 holds, then the OLS estimator of β obtained by using only the observations with complete data on all covariates is unbiased for β .*

Proof. The OLS estimator for the complete data may be written

$$\hat{\beta} = (\mathbf{X}'\mathbf{D}\mathbf{X})^{-1}\mathbf{X}'\mathbf{D}\mathbf{y},$$

where \mathbf{D} is an $N \times N$ diagonal matrix whose n -th diagonal element d_n takes value 1 if no covariate is missing for the n -th observation and value 0 otherwise. The elements of \mathbf{D} are related to the elements of the missing-indicator matrix \mathbf{M} through $d_n = \prod_{k=1}^K (1 - m_{nk})$. The Ignorability Assumption implies that any function of \mathbf{M} , in particular \mathbf{D} , is independent of \mathbf{y} conditional on \mathbf{X} . From (2),

$$\mathbb{E}(\hat{\beta} \mid \mathbf{X}, \mathbf{D}) = (\mathbf{X}'\mathbf{D}\mathbf{X})^{-1}\mathbf{X}'\mathbf{D}\mathbf{X}\beta = \beta,$$

and therefore $\mathbb{E}(\hat{\beta}) = \beta$. \square

An implication of Theorem 1 is that the subsample with complete data satisfies

$$\mathbf{y}^0 = \mathbf{X}^0\beta + \mathbf{u}^0, \quad (4)$$

where \mathbf{u}^0 is an $N_0 \times 1$ vector of homoskedastic and serially uncorrelated regression errors. This result supports the common practice of complete-case analysis, namely estimating β by regressing \mathbf{y}^0 on \mathbf{X}^0 . However, severe loss of information, and hence of precision, may result unless the fraction of deleted cases is small.

3. THE AUGMENTED MODEL WITH AUXILIARY VARIABLES

Suppose that, for each subsample $j = 1, \dots, J$ with incomplete data, the values of the K_j^* missing covariates are filled-in using some imputation procedure. A covariate with imputed values is called an imputed covariate. The $N_j \times K_j^*$ matrix corresponding to the set of imputed covariates is called the imputation matrix for the j -th subsample and is denoted by \mathbf{L}^j . The $N_j \times K$ matrix $\mathbf{W}^j = [\mathbf{X}_a^j, \mathbf{L}^j]$, whose columns correspond to the K_j available covariates and the K_j^* imputed covariates, is called the completed design matrix for the j -th subsample. Our treatment of imputation is very general and covers a variety of imputation procedures, including regression and donor-based methods such as nearest-neighbor and hot-deck imputations. It also allows for the possibility that different imputation procedures are used for different covariates, or for different subsamples with incomplete data.

Consider modelling the $N_j \times 1$ outcome vector \mathbf{y}^j for the j -th subsample as a linear function of the observed covariates in \mathbf{W}^j . The best (minimum mean-square error) linear predictor of \mathbf{y}^j given $\mathbf{W}^j = [\mathbf{X}_a^j, \mathbf{L}^j]$ is

$$\begin{aligned} \mathbb{E}^*(\mathbf{y}^j \mid \mathbf{X}_a^j, \mathbf{L}^j) &= \mathbb{E}^*(\mathbf{X}^j\beta \mid \mathbf{X}_a^j, \mathbf{L}^j) \\ &= \mathbf{X}_a^j\beta_a^j + \mathbb{E}^*(\mathbf{X}_m^j \mid \mathbf{X}_a^j, \mathbf{L}^j)\beta_m^j \\ &= \mathbf{X}_a^j\beta_a^j + (\mathbf{X}_a^j\Delta^j + \mathbf{L}^j\Gamma^j)\beta_m^j \\ &= \mathbf{X}_a^j\gamma_a^j + \mathbf{L}^j\gamma_m^j, \end{aligned}$$

where β_a^j and β_m^j are the subvectors of β associated with \mathbf{X}_a^j and \mathbf{X}_m^j respectively, $\mathbb{E}^*(\mathbf{X}_m^j \mid \mathbf{X}_a^j, \mathbf{L}^j) = \mathbf{X}_a^j\Delta^j + \mathbf{L}^j\Gamma^j$ is the best linear predictor of \mathbf{X}_m^j given \mathbf{X}_a^j and \mathbf{L}^j , and

$$\gamma_a^j = \beta_a^j + \Delta^j\beta_m^j, \quad \gamma_m^j = \Gamma^j\beta_m^j.$$

The resulting linear model for the j -th subsample may be written, more compactly,

$$\mathbf{y}^j = \mathbf{W}^j \boldsymbol{\gamma}^j + \mathbf{v}^j, \quad j = 1, \dots, J, \quad (5)$$

where $\boldsymbol{\gamma}^j$ is the $K \times 1$ vector consisting of the coefficients associated with the observed and the imputed covariates in $\mathbf{W}^j = [\mathbf{X}_a^j, \mathbf{L}^j]$, and \mathbf{v}^j is an $N_j \times 1$ vector of projection errors that, by construction, has mean zero and is orthogonal to \mathbf{W}^j .

Two important features distinguish model (5) from the original model (1). First, the vector of population coefficients $\boldsymbol{\gamma}^j$ is generally different from $\boldsymbol{\beta}$ unless $\Delta^j = 0$ and Γ^j is equal to the identity matrix or, equivalently, $\mathbb{E}^*(\mathbf{X}_m^j | \mathbf{X}_a^j, \mathbf{L}^j) = \mathbf{L}^j$, that is, given the imputations, the available covariates contain no further information about the missing covariates. Second, the elements of the error vector \mathbf{v}^j are not necessarily homoskedastic, even when homoskedasticity holds for the corresponding subvector \mathbf{u}^j of \mathbf{u} .

Letting $\boldsymbol{\delta}^j = \boldsymbol{\gamma}^j - \boldsymbol{\beta}$, $j = 1, \dots, J$, and stacking on top of each other the complete-case model and the J linear models for the subsamples with incomplete data gives

$$\begin{bmatrix} \mathbf{y}^0 \\ \mathbf{y}^* \end{bmatrix} = \begin{bmatrix} \mathbf{X}^0 \\ \mathbf{W}^* \end{bmatrix} \boldsymbol{\beta} + \begin{bmatrix} \mathbf{0} \\ \mathbf{Z}^* \end{bmatrix} \boldsymbol{\delta} + \begin{bmatrix} \mathbf{u}^0 \\ \mathbf{v}^* \end{bmatrix},$$

where

$$\mathbf{y}^* = \begin{bmatrix} \mathbf{y}^1 \\ \vdots \\ \mathbf{y}^J \end{bmatrix}, \quad \mathbf{W}^* = \begin{bmatrix} \mathbf{W}^1 \\ \vdots \\ \mathbf{W}^J \end{bmatrix}, \quad \mathbf{Z}^* = \begin{bmatrix} \mathbf{W}^1 & & \\ & \ddots & \\ & & \mathbf{W}^J \end{bmatrix}, \quad \mathbf{v}^* = \begin{bmatrix} \mathbf{v}^1 \\ \vdots \\ \mathbf{v}^J \end{bmatrix},$$

and $\boldsymbol{\delta}$ is the $JK \times 1$ vector consisting of $\boldsymbol{\delta}^1, \dots, \boldsymbol{\delta}^J$. We can now write the model for the available and the imputed data as the grand model

$$\mathbf{y} = \mathbf{W} \boldsymbol{\beta} + \mathbf{Z} \boldsymbol{\delta} + \mathbf{v}, \quad (6)$$

where $\boldsymbol{\beta}$ is the parameter of primary interest, $\boldsymbol{\delta}$ is a vector of nuisance parameters, and

$$\mathbf{y} = \begin{bmatrix} \mathbf{y}^0 \\ \mathbf{y}^* \end{bmatrix}, \quad \mathbf{W} = \begin{bmatrix} \mathbf{X}^0 \\ \mathbf{W}^* \end{bmatrix}, \quad \mathbf{Z} = \begin{bmatrix} \mathbf{0} \\ \mathbf{Z}^* \end{bmatrix}, \quad \mathbf{v} = \begin{bmatrix} \mathbf{u}^0 \\ \mathbf{v}^* \end{bmatrix},$$

respectively an N -vector, an $N \times K$ matrix, an $N \times JK$ matrix, and an N -vector of errors that has mean zero and is orthogonal to both \mathbf{W} and \mathbf{Z} . Notice that the matrix \mathbf{W} is obtained by filling-in the missing covariate values with the available imputations. Model (6) includes all observations re-ordered groupwise: first the complete cases, then the first group with incomplete data, etc. Ordering of the groups is arbitrary and plays no role in the analysis. In the terminology of Danilov and Magnus [4], the K columns of \mathbf{W} are the “focus” regressors, while the JK columns of \mathbf{Z} are the “auxiliary” regressors.

4. A MISSING-INDICATOR INTERPRETATION

Before presenting our main result it is instructive to give a missing-indicator interpretation of model (6). Indeed, the JK auxiliary variables in the matrix \mathbf{Z} are obtained by multiplying the covariates in each group by the various indicators of group membership. To see this write $\mathbf{Z} = [\mathbf{Z}_1, \dots, \mathbf{Z}_j, \dots, \mathbf{Z}_J]$, where \mathbf{Z}_j is the $N \times K$ matrix that contains the auxiliary variables for the j -th group. Let $\mathbf{1}_K$ denote the $1 \times K$ vector whose elements are all equal to one and let \mathbf{d}_j denote the $N \times 1$ vector of group-membership indicators

for the j -th group (the elements of \mathbf{d}_j are equal to one for observations in group j and zero otherwise). Then

$$\mathbf{Z}_j = [\mathbf{1}_K \otimes \mathbf{d}_j] \cdot \mathbf{W}, \quad j = 1, \dots, J,$$

where \otimes denotes the Kronecker product and \cdot the Hadamard (elementwise) product.

As an illustration, consider the linear model

$$\mathbb{E}(y_n \mid x_{n1}, x_{n2}) = \beta_0 + \beta_1 x_{n1} + \beta_2 x_{n2},$$

with a constant term and two covariates, x_1 and x_2 . Suppose that, in addition to the group with complete data, one has two groups with incomplete data: in group 1 [resp. 2] only the first [resp. second] covariate is missing. If d_{n0} , d_{n1} and d_{n2} denote the group-membership indicators, and L_{n1}^1 and L_{n2}^2 denote the imputed values in each group with incomplete data, then we may write

$$\begin{aligned} y_n = & d_{n0}(\beta_0 + \beta_1 x_{n1} + \beta_2 x_{n2}) + d_{n1}(\gamma_0^1 + \gamma_1^1 L_{n1}^1 + \gamma_2^1 x_{n2}) + \\ & + d_{n2}(\gamma_0^2 + \gamma_1^2 x_{n1} + \gamma_2^2 L_{n2}^2) + v_n. \end{aligned}$$

Let w_{nk} be equal to x_{nk} if the k -th covariate is observed for the n -th observation and to its imputed value otherwise. Then the last relation may be written

$$\begin{aligned} y_n = & \beta_0 + \beta_1 w_{n1} + \beta_2 w_{n2} + \delta_0^1 d_{n1} + \delta_0^2 d_{n2} + \\ & + \delta_1^1 d_{n1} w_{n1} + \delta_1^2 d_{n2} w_{n1} + \delta_2^1 d_{n1} w_{n2} + \delta_2^2 d_{n2} w_{n2} + v_n, \end{aligned}$$

where $\delta_k^j = \gamma_k^j - \beta_k$. This is exactly model (6) for this special case, where the auxiliary variables added to the w_k 's are the group-membership indicators and their interactions with the constant term and the observed or imputed covariates.

5. MAIN RESULT

The following result shows that, no matter which imputation procedure is chosen, the OLS estimate of β in the grand model (6) and in the complete-case model (4) are numerically the same. Thus, the statistical properties of the two estimators are also the same. In particular, if the latter is unbiased (for example, the conditions of Theorem 1 hold), so is the former.

Theorem 2. *Suppose that the matrix \mathbf{W} is of full column rank K . Then, for any choice of imputation matrices $\mathbf{L}^1, \dots, \mathbf{L}^J$, the OLS estimate of β in the “grand” model (6) coincides with the OLS estimate of β in the complete-case model (4).*

Proof. Given any matrix \mathbf{A} , let $\mathbf{R}_\mathbf{A} = \mathbf{I} - \mathbf{A}(\mathbf{A}'\mathbf{A})^- \mathbf{A}'$, where $(\mathbf{A}'\mathbf{A})^-$ denotes a g -inverse of $\mathbf{A}'\mathbf{A}$. Since $\mathbf{Z}'\mathbf{Z}$ and \mathbf{Z} have the same rank, the rank of $\mathbf{Z}(\mathbf{Z}'\mathbf{Z})^- \mathbf{Z}'$ is equal to the rank of \mathbf{Z} (Rao and Mitra [20]). The fact that the rank of \mathbf{Z} may be less than JK , implies that the rank of $\mathbf{R}_\mathbf{Z}$ must be at least $N - JK$. Because $N \geq K(J + 1)$ implies that $K \leq N - JK$, it follows that the rank of \mathbf{W} cannot exceed the rank of $\mathbf{R}_\mathbf{Z}$, so the matrix $\mathbf{W}'\mathbf{R}_\mathbf{Z}\mathbf{W}$ must be nonsingular. Thus, by the Frisch-Waugh-Lovell (Partitioned Regression) Theorem, the OLS estimate of β in model (6) is

$$\hat{\beta} = (\mathbf{W}'\mathbf{R}_\mathbf{Z}\mathbf{W})^{-1} \mathbf{W}'\mathbf{R}_\mathbf{Z}\mathbf{y} = (\tilde{\mathbf{X}}'\tilde{\mathbf{X}})^{-1} \tilde{\mathbf{X}}'\mathbf{y},$$

where $\tilde{\mathbf{X}} = \mathbf{R}_Z \mathbf{W}$. Next notice that

$$\tilde{\mathbf{X}} = \begin{bmatrix} \mathbf{I}_{N_0} & & & \\ & \mathbf{R}_{\mathbf{W}^1} & & \\ & & \ddots & \\ & & & \mathbf{R}_{\mathbf{W}^J} \end{bmatrix} \begin{bmatrix} \mathbf{X}^0 \\ \mathbf{W}^1 \\ \vdots \\ \mathbf{W}^J \end{bmatrix} = \begin{bmatrix} \mathbf{X}^0 \\ \mathbf{0} \\ \vdots \\ \mathbf{0} \end{bmatrix},$$

where we used the fact that $\mathbf{W}^j(\mathbf{W}^{j'}\mathbf{W}^j)^{-}\mathbf{W}^{j'}\mathbf{W}^j = \mathbf{W}^j$ for all j and any choice of g -inverse (Rao and Mitra [20]). Therefore $\tilde{\mathbf{X}}'\tilde{\mathbf{X}} = \mathbf{X}^{0'}\mathbf{X}^0$ and $\tilde{\mathbf{X}}\mathbf{y} = \mathbf{X}^{0'}\mathbf{y}^0$. Hence

$$\hat{\boldsymbol{\beta}} = (\tilde{\mathbf{X}}'\tilde{\mathbf{X}})^{-1}\tilde{\mathbf{X}}'\mathbf{y} = (\mathbf{X}^{0'}\mathbf{X}^0)^{-1}\mathbf{X}^{0'}\mathbf{y}^0,$$

which is the complete-case estimate. \square

The matrix \mathbf{W} is of full column rank if, as we already assumed, the K columns of \mathbf{X}^0 are linearly independent. The use of a g -inverse in the proof of the theorem is necessary because some of the completed design matrices \mathbf{W}^j may be singular, which happens if $N_j < K$ or if $N_j \geq K$ but the columns of \mathbf{W}^j are linearly dependent. The latter is for example the case when a missing covariate value is replaced by its average value computed from the available cases (mean imputation) or by its predicted values based on the observed covariates \mathbf{X}_a^j and the coefficients from an OLS regression using the subsample with complete data (deterministic regression imputation). One can replace a g -inverse with the regular inverse when the J subsamples with incomplete data are such that all \mathbf{W}^j 's have full column rank. In practice, this may be achieved by dropping groups that contain too few observations and avoiding mean imputation or deterministic regression imputation. After all, these two imputation methods are known to produce completed data sets with undesirable properties, for example they have less variability than a set of truly observed values (see e.g. Lundström and Särndal [14]).

The complete-case model (4) and the grand model (6) may at first appear as two polar approaches to the problem of handling missing data in model (1). At one extreme is complete-case analysis. Under the assumption of Theorem 1, this gives an unbiased estimate of $\boldsymbol{\beta}$ but may throw away too much information by retaining only the observations in the subsample with complete data. At the other extreme, all observations are retained but some imputation procedure is adopted to fill-in the missing data. In fact, Theorem 2 shows that if $\boldsymbol{\beta}$ and $\boldsymbol{\delta}$ in (6) are left unconstrained then this second approach is equivalent to complete-case analysis as far as estimation of $\boldsymbol{\beta}$ is concerned.

A referee offered the following heuristic. Our model places no restrictions (equivalently, uses no information) on the imputation method. In the decomposition $\boldsymbol{\gamma}^j = \boldsymbol{\beta} + \boldsymbol{\delta}^j$, it is only the observations from the complete case that sort out the part that should be $\boldsymbol{\beta}$. Since the remaining cases provide absolutely no information, the estimates are the same. “No information added, no change”.

The standard practice of regressing \mathbf{y} only on the completed design matrix \mathbf{W} omitting the variable in \mathbf{Z} corresponds to using a restricted version of the grand model (6) where all elements of the vector $\boldsymbol{\delta}$ are set equal to zero. This is the same as assuming that the missing data mechanism satisfies Assumption 1 and the imputation procedure is such that $\boldsymbol{\beta} = \boldsymbol{\gamma}^j$ for each $j = 1, \dots, J$. The less frequent practice of regressing \mathbf{y} on \mathbf{W} and

the set of group-membership indicators (which we shall refer to as the simple missing-indicator method) corresponds to using another restricted version of the grand model, where all interactions between the group membership indicators and the observed or derived covariates are set equal to zero.

Both sets of restrictions are testable. Testing the first set of restrictions corresponds to testing the hypothesis that all regression coefficients do not change across the J groups containing missing covariates, while testing the second set of restrictions corresponds to testing the hypothesis that, except for the intercepts, all other regression coefficients do not change across the J groups containing missing covariates.

The precise nature of these tests, in particular the form of the test statistics, depends on the properties of the error vector \mathbf{v} in model (6). Given OLS estimates $\hat{\beta}$ and $\hat{\delta}$ of β and δ in the grand model, classical F -tests would be appropriate if it can safely be assumed that \mathbf{v} is a vector of homoskedastic and serially uncorrelated regression errors. If this assumption cannot be justified, then one could use a “robustified” version of these tests based on an estimator of the sampling variance of $\hat{\beta}$ and $\hat{\delta}$ that is consistent under heteroskedasticity or autocorrelation of unknown form in the elements of \mathbf{v} .

In our view, however, the key issue is not what statistic to use for testing, but whether it makes sense to ask questions such as: Is it true that $\delta = 0$? Following Leamer [11] and Magnus and Durbin [15], we think that asking such questions in this context is wrong. The right question is: What is the best available estimator of β ?

6. BIAS VERSUS PRECISION

Theorem 2 says that unbiased estimates of β may be obtained in two equivalent ways, either by using the N_0 observations in the subsample with complete data, or by using all N observations and the grand model (6) which includes the imputed values of the missing covariates in the matrix \mathbf{W} and the auxiliary variables in the matrix \mathbf{Z} . We also know from standard results that placing restrictions on the elements of δ may lead to biased but more precise estimates of β .

Two approaches may be followed to handle this trade-off between bias and precision in the estimation of β : model reduction and model averaging. Either approach can be applied to model (6).

Model reduction involves first selecting an intermediate model between the grand model and the fully restricted model where $\delta = 0$, and then estimating the parameter of interest β conditional on the selected model. Model reduction may be carried out through variable selection methods, such as stepwise regression (see e.g. Kennedy [10]), or more complex General-to-Specific procedures (see Campos, Ericsson and Hendry [1] for a survey). The details of the model reduction procedure may also depend on whether one allows dropping arbitrary subsets of auxiliary variables in \mathbf{Z} , or only subsets of auxiliary variables corresponding to specific subsamples with missing covariates. Dropping one of the columns of \mathbf{Z} amounts to selecting a group j and, in the corresponding equation (5), restricting one element of δ^j to zero. This in turn corresponds to forcing the coefficient of that particular covariate in the completed design matrix \mathbf{W}^j to be the same as in the subsample with complete data. Dropping the columns of \mathbf{Z} corresponding to the j -th subsample amounts instead to restricting all element of δ^j to be zero, which in turn

corresponds to forcing the relationship between \mathbf{y}^j and the completed design matrix \mathbf{W}^j to be the same as that between \mathbf{y}^0 and \mathbf{X}^0 in the subsample with complete data.

One well known problem with this approach is pre-testing. Another is that model selection and estimation are completely separated. As a result, the reported conditional estimates tend to be interpreted as if they were unconditional.

Model averaging is different. Instead of selecting a model out of the available set of models, one first estimates the parameter of interest β conditional on each model in the model space, and then computes the estimator of β as a weighted average of these conditional estimators. When the model space contains I models, a model averaging estimator of β is of the form

$$\bar{\beta} = \sum_{i=1}^I \lambda_i \hat{\beta}_i, \quad (7)$$

where the λ_i are non-negative weights that add up to one, and $\hat{\beta}_i$ is the estimator of β obtained by conditioning on the i th model. In our case, the model space contains 2^{JK} models, a huge number unless both J and K are small.

In Bayesian model averaging (BMA), each $\hat{\beta}_i$ is weighted by the posterior probability of the corresponding model. If equal prior probabilities are assigned to each model under consideration, then the λ_i are just proportional to the marginal likelihood of each model. Bayesian averaging of both classical (least-squares) and Bayesian estimators have been considered, with the posterior mean of β for the model under consideration as the typical Bayesian estimator. Bayesian averaging of Bayesian estimators has been popularized by Raftery, Madigan and Hoeting [19], while Bayesian averaging of classical estimators has been popularized by Sala-i-Martin, Doppelhofer and Miller [22]. The choice between the different approaches involves considering the computational burden and the statistical properties of the resulting estimators and, in the case of BMA, the nature of the assumed priors. The role of priors would also arise if a Bayesian model reduction approach is taken.

Magnus, Powell and Prüfer [16] study the properties of model averaging estimators of the same form as (7) with $\lambda_i = \lambda_i(\tilde{\mathbf{v}})$, where $\tilde{\mathbf{v}}$ is the vector of OLS residuals from the regression of \mathbf{y} on \mathbf{W} only. Their class of weighted-average least squares (WALS) estimators generalizes to the case when $I \geq 2$ the class of estimators introduced by Magnus and Durbin [15], which contains the classical pre-test estimator as a special case. Although WALS estimators are in fact BMA estimators, they differ from standard BMA in three important respects: their computational burden, the choice of prior for δ , and their statistical properties.

The main advantage of WALS is that, although we have 2^{JK} models, the computational burden is only proportional to JK . With medium or large values of J or K , the computation burden is minimal compared to standard BMA.

Like standard BMA, WALS assume a classical Gaussian linear model for (6) and noninformative priors for β and the error variance σ^2 . The assumption that the regression errors are homoskedastic and serially uncorrelated is not crucial for WALS, and the method can be generalized to non-spherical errors (Magnus, Wan and Zhang [17]). The key step in WALS is to reparameterize the model replacing $\mathbf{Z}\delta$ by $\mathbf{Z}^*\delta^*$, with $\mathbf{Z}^* = \mathbf{Z}\mathbf{P}\mathbf{\Lambda}^{-1/2}$ and $\delta^* = \mathbf{\Lambda}^{1/2}\mathbf{P}'\delta$, where \mathbf{P} is an orthonormal matrix and $\mathbf{\Lambda}$ is a

diagonal matrix such that $\mathbf{P}'\mathbf{Z}'\mathbf{R}_w\mathbf{Z}\mathbf{P} = \mathbf{\Lambda}$. The main difference with respect to standard BMA is that, instead of a multivariate Gaussian prior for $\boldsymbol{\delta}$, WALs use a Laplace distribution with zero mean for the independently and identically distributed elements of the transformed parameter vector $\boldsymbol{\eta} = \boldsymbol{\delta}^*/\sigma$, whose i th element, η_i is the population t -ratio on δ_i , the i th element of $\boldsymbol{\delta}$. In this formulation, ignorance is a situation where it is equally likely for these population t -ratios to be larger or smaller than one in absolute value.

Finally, unlike standard BMA, WALs have bounded risk and are near-optimal in terms of a well-defined regret criterion (Magnus, Powell and Prüfer [16]).

7. AN APPLICATION

In this section we apply our approach in the context of a concrete example with missing data. The problem at hand is that of estimating the relation between body-mass and income using survey data affected by item non-response. We first present the estimates one obtains by the complete-case approach, by using raw data (no dummies), and by the simple indicator method. We then compare them with the estimates one obtains using different model-selection or model-averaging techniques on the basis of the grand model (6).

The body-mass index (BMI), namely the ratio of weight (in kg) to squared height (in meters), is one way of combining weight and height into a single measure. Due to its ease of measurement and calculation, the BMI is the most common diagnostic tool to identify obesity problems within a population. As such, it has received lots of attention in the recent literature on the obesity epidemic and its economic and public health consequences (Cutler, Glaeser and Shapiro [3], Philipson and Posner [18]).

The obesity epidemic is essentially an imbalance between food intake and energy expenditure. It has been argued that this imbalance may be linked to income (see e.g. Drewnowski and Specter [5]). The available empirical evidence (Cawley, Moran and Simon [2]) for elderly people in the USA, Sanz-de-Galdeano [23] and García-Villar and Quintana-Domeque [6] for Europe) is inconclusive for men, whereas for women there appears to be a more clear indication of a negative correlation between BMI and income.

Our data are from Release 2 of the 1st wave of the Survey of Health, Ageing, and Retirement in Europe (SHARE), a multidisciplinary and cross-national household panel survey designed to investigate several aspects of the elderly population in Europe. The target population of SHARE consists of people aged 50+ living in residential households, plus their co-resident partners irrespective of age. The 1st wave, conducted in 2004, covered 15,544 households and 22,431 individuals in 11 European countries. All national samples are selected through probability sampling.

The key to ensure comparability is the adoption of a common survey instrument. The Physical Health module of the questionnaire collects self-reported height and weight, the Income module collects information on 25 income components which are then aggregated into a measure of household income, and the Consumption module collects household expenditure on four consumption categories (food at home, food outside the home, telephone, all goods and services) in the last month.

Nonresponse to household income and food expenditure is substantial, and in this case we use the imputations provided by SHARE. Complete or partial nonresponse

to household income occurs for as much as 60 percent of the observations, such a high fraction being due to the fact that this variable is obtained by aggregating a large number of income components across household members. Nonresponse to food expenditure occurs for about 15 percent of the observations.

To impute missing values, SHARE uses a complex two-stage multivariate procedure (Kalwij and van Soest [9]). Imputations are first obtained recursively for a few core variables. In the second stage, the imputed values from the first stage are used to impute the other variables. This procedure essentially employs only univariate regression imputation methods. It is important to notice that height and weight are never used to impute missing variables. To allow multiple imputation methods, SHARE provides 5 imputations for each missing value. SHARE imputes total household income by separately imputing each income component and then aggregating them. Imputations are provided for individual incomes of all eligible partners who did not agree to participate to the survey.

We focus on the income-BMI relationship for males. We model the mean of log BMI as a function of age and age squared, log household income per capita, log household food expenditure per capita, and a dummy indicator for low educational attainment. In addition to the subsample with complete data (4,067 obs., 35.5%), we have three subsamples with incomplete data: one where only food expenditure is missing (287 obs., 2.5%), one where only household income is missing (5,891 obs., 51.3%), and one where both household income and food expenditure are missing (1,230 obs., 10.7%). For each variable, we use the first of the 5 available imputations.

Table 1 shows the estimated coefficients for age and its square (*agesq*), log household income per capita (*lypc*), log food expenditure per capita (*lfpc*), and a dummy for not having a high-school degree (*lowed*). The first three columns contain estimates for the complete-case/grand model, the naive estimator corresponding to $\delta = 0$, and the simple missing-indicator method. The other three are obtained by model-selection or model-averaging on the basis of the grand model (6): Stata's backward-stepwise selection with p -value equal to .05, WALS and BMA. The standard errors in the table have all been obtained using the "wild" bootstrap with 50 replications and are robust to heteroskedasticity of the errors in the grand model.

The BMA and WALS estimates have been computed using the Matlab code downloaded from Jan Magnus's web page at <http://center.uvt.nl/staff/magnus/wals/>. This BMA implementation estimates all possible models, so it becomes very time consuming when J or K are large. In our case, with $J = 3$ subsamples with incomplete data and $K = 6$ focus regressors (including the constant term), examining all possible 2^{18} models required about one day on our desktop computer. Faster implementations are available, but they estimate only a randomly chosen subset of all possible models and have the important disadvantage of not using the distinction between focus and auxiliary regressors, which is key to our analysis.

As for WALS, it is worth discussing briefly the concept and treatment of uncertainty implicit in the choice of a Laplace prior for the elements of the transformed parameter vector η . Assuming this particular prior means that we think that it is equally likely that the observed value of the t -statistic on any element of δ is greater or smaller than one. This is equivalent to say that we are agnostic about the quality of the imputation: it

could be either good or bad. Since we are simply users, not producers, of the imputations, this may not be a bad assumption.

There is agreement between the different methods on the qualitative effect of the various variables: concave for age, negative for education and income, and positive but not statistically significant for food expenditure. The magnitude of the estimated coefficients, however, differs considerably across methods. At one extreme are the naive estimator, the simple missing-indicator method and BMA, that produce nearly identical results: they assign more importance to age and less importance to income. At the other extreme are the complete-case estimator and WALs: they assign less importance to age and more importance to income. It is noteworthy that, in this example, WALs is close to complete-case (all dummies in the model), while BMA is close to naive estimates (no dummy). Thus, starting with the the grand model, WALs seems to give more weight to the auxiliary dummies than BMA. The stepwise procedure gives estimates of the relative effects of age and income that are somewhat in between these two extremes.

8. CONCLUDING REMARKS

In this paper we formalized the trade-off between bias and efficiency that arises when there are missing covariate values in a regression relationship of interest and showed how to tackle this trade-off by model reduction procedures or model averaging methods. In future work we plan to extend our approach to generalized linear models (GLM), such as logit, probit and Poisson regression, for which we conjecture that versions of Theorems 1 and 2 also hold. Our conjecture is motivated by the fact that maximum-likelihood estimators of exponential family models may be obtained by iteratively reweighted least squares.

REFERENCES

- [1] Campos, J., N.R. Ericsson and D.F. Hendry (2005): General-to-Specific Modeling: An Overview and Selected Bibliography, FRB International Finance Discussion Paper No. 838. Available at SSRN: <http://ssrn.com/abstract=791684>
- [2] Cawley, J., J.R. Moran and K.I. Simon (2008): The Impact of Income on the Weight of Elderly Americans, NBER Working Paper No. 14104.
- [3] Cutler, D.A., E.L. Glaeser and J.M. Shapiro (2003): Why Have Americans Become More Obese?, *Journal of Economic Perspectives* **17**, pp. 93–118.
- [4] Danilov, D., and J.R. Magnus (2004): On the Harm That Ignoring Pretesting Can Cause, *Journal of Econometrics* **122**, pp. 27–46.
- [5] Drewnowski, A., and S. Specter (2004): Poverty and Obesity: The Role of Energy Density and Costs, *American Journal of Clinical Nutrition* **79**, pp. 6–16.
- [6] García Villar, J., and C. Quintana-Domeque (2009): Income and Body Mass Index in Europe, *Economics & Human Biology* **7**, pp. 73–83.
- [7] Horton, N.J., and K.P. Kleinman (2007): Much Ado About Nothing: A Comparison of Missing Data Methods and Software to Fit Incomplete Data Regression Models, *The American Statistician* **61**, pp. 79–90.
- [8] Jones, M.P. (1996): Indicator and Stratification Methods for Missing Explanatory Variables in Multiple Linear Regression, *Journal of the American Statistical Association* **91**, pp. 222–230.
- [9] Kalwij, A., and A. van Soest (2005): Item Non-Response and Alternative Imputation Procedures, in Axel Börsch-Supan and Hendrik Jürges (eds.), *The Survey of Health, Aging, and Retirement in Europe—Methodology*, MEA, Mannheim, pp. 128–150.
- [10] Kennedy Jr., W.J., and T.A. Bancroft (1971): Model-Building for Prediction in Regression Based on Repeated Significance Tests, *Annals of Mathematical Statistics* **42**, pp. 1273–1284.
- [11] Leamer, E.E. (1978): *Specification Searches. Ad Hoc Inference with Nonexperimental Data*, Wiley.
- [12] Little, R.J.A. (1992): Regression with Missing X's: A Review, *Journal of the American Statistical Association* **87**, pp. 1227–1237.
- [13] Little, R.J.A., and D. Rubin (2002): *Statistical Analysis with Missing Data* (2nd ed.), Wiley.
- [14] Lundström, S., and C.-E. Särndal (2002): *Estimation in the Presence of Nonresponse and Frame Imperfections*, Statistics Sweden.
- [15] Magnus, J.R., and J. Durbin (1999): Estimation of Regression Coefficients of Interest When Other Regression Coefficients are of no Interest, *Econometrica* **67**, pp. 639–643.
- [16] Magnus, J.R., O. Powell and P. Prüfer (2010): A Comparison of Two Averaging Techniques With an Application to Growth Empirics, *Journal of Econometrics* **154**, pp. 139–153.
- [17] Magnus, J.R., A.T.K. Wan and X. Zhang (2010): WALS Estimation with Nonspherical Disturbances and an Application to the Hong Kong Housing Market, *Computational Statistics & Data Analysis*, forthcoming.
- [18] Philipson, T., and R. Posner (2008): Is the Obesity Epidemic a Public Health Problem? A Decade of Research on the Economics of Obesity, NBER Working Paper No. 14010.
- [19] Raftery, A.E., D. Madigan and J.A. Hoeting (1997): Bayesian Model Averaging for Linear Regression Models, *Journal of the American Statistical Association* **92**, pp. 179–191.
- [20] Rao, C.R., and S.L. Mitra, D. (1971): *Generalized Inverse of Matrices and its Applications*, Wiley.
- [21] Rubin, D. (1987): *Multiple Imputations*, Wiley.
- [22] Sala-i-Martin, X., G. Doppelhofer, and R.I. Miller (2004): Determinants of Long-Term Growth: A Bayesian Averaging of Classical Estimates (BACE) Approach, *American Economic Review* **94**, pp. 813–835.
- [23] Sanz-de-Galdeano, A. (2005): The Obesity Epidemic in Europe, IZA Discussion Paper No. 1814.
- [24] Wooldridge, J.M. (2002): *Econometric Analysis of Cross Section and Panel Data*, MIT Press.

TABLE 1. Estimated coefficients (bootstrap standard errors in parentheses).

	Complete case	Naive	Simple indicator	Stata's stepwise	WALS	BMA
age	0.0008 (0.0006)	0.0023 *** (0.0005)	0.0023 *** (0.0005)	0.0020 *** (0.0007)	0.0012 ** (0.0005)	0.0023 *** (0.0005)
agesq	-0.0082 *** (0.0016)	-0.0107 *** (0.0012)	-0.0107 *** (0.0012)	-0.0108 *** (0.0018)	-0.0087 *** (0.0015)	-0.0107 *** (0.0012)
lypc	-0.0196 *** (0.0028)	-0.0097 *** (0.0015)	-0.0098 *** (0.0015)	-0.0174 *** (0.0034)	-0.0165 *** (0.0023)	-0.0101 *** (0.0015)
lfpc	0.0054 (0.0049)	0.0042 (0.0028)	0.0044 (0.0028)	0.0049 (0.0034)	0.0046 (0.0043)	0.0042 (0.0028)
lowed	0.0144 *** (0.0046)	0.0201 *** (0.0030)	0.0201 *** (0.0030)	0.0197 *** (0.0047)	0.0160 *** (0.0041)	0.0201 *** (0.0030)
No. obs.	4,067	11,475	11,475	11,475	11,475	11,475

Note: Observed p -values: * $p < .10$; ** $p < .05$; *** $p < .01$.