sociological science

Supplement to:

Winship, Christopher, and Bruce Western. 2016. "Multicollinearity and Model Misspecification." Sociological Science 3: 627-649.

## Appendix A. Deriving the Posterior Mean and Variance of $\mathbf{b}$

The approach here follows Leamer (1991) and Western (1999). Start with the basic model equation:

$$\mathbf{y} = \mathbf{X}\mathbf{b} + \mathbf{e} \tag{A.1}$$

Consider an auxiliary equation indicating a possible linear dependence of the error term, $\mathbf{e}$ on $\mathbf{X}$:

$$\mathbf{e} = \mathbf{X}\mathbf{a} + \boldsymbol{\varepsilon} \tag{A.2}$$

where $\mathbf{a}$ is a $\boldsymbol{k} \times \mathbf{1}$. Here $\mathbf{a} = (\mathbf{X'X})^{-1}\mathbf{X'e}$, which is equal to the bias factor for $\mathbf{b}_{\boldsymbol{OLS}}$ in equation 5. Substituting (A.2) into (A.1) yields:

$$\mathbf{y} = \mathbf{X}\mathbf{b} + \mathbf{X}\mathbf{a} + \boldsymbol{\varepsilon} \tag{A.3}$$

In equation (A.3) $\mathbf{b}$ and $\mathbf{a}$ cannot be estimated from the data alone. Leamer (1991) and Zellner (1971) observed, however, that equations similar to (A3) can be estimated with Bayesian methods as long as some coefficients are given proper prior distributions. The intuition here is that priors for $\mathbf{b}$ and $\mathbf{a}$ contain distinct and separate information on their possible values. The data (possibly along with prior information) allows one to estimate $(\mathbf{b} + \mathbf{a})$ and prior information allows one to separate this value into separate estimates for $\mathbf{b}$ and $\mathbf{a}$. Classic OLS is an extreme version of this in that it simply assumes that $\mathbf{a} = \mathbf{0}$.

We now derive posterior means and variances for $\mathbf{b}$ and $\mathbf{a}$. Rewrite (A.3) as

$$\mathbf{y} = \mathbf{Z}\boldsymbol{\beta} + \boldsymbol{\varepsilon},$$

where $\mathbf{Z} = [\mathbf{X}\ \mathbf{X}]$ and $\boldsymbol{\beta}' = [\mathbf{b}'\ \mathbf{a}']$, and $\boldsymbol{\varepsilon}$ is assumed to be iid normal. The coefficient vector, $\boldsymbol{\beta}$, is assumed to have a normal prior with mean, $\boldsymbol{\beta}'_{\mathbf{0}} = [\mathbf{b}'_{\mathbf{0}}\ \mathbf{a}'_{\mathbf{0}}]$ and variance,

$$\mathbf{V_0} = \begin{bmatrix} \mathbf{P_0} & \mathbf{0} \\ \mathbf{0} & \mathbf{Q_0} \end{bmatrix},$$

where $\mathbf{P_0}$ is the prior covariance matrix for $\mathbf{b}$ and $\mathbf{Q_0}$ is the prior covariance matrix for $\mathbf{a}$. A prior with this structure assumes that beliefs about the coefficients $\mathbf{b}$ are independent of beliefs about the model misspecification.

Closed-form expressions are available for the posterior mean and variance of $\boldsymbol{\beta}$ where the coefficients have a normal prior (Gelman et al. 1995). The

1

posterior variance is given by:

$$\begin{aligned}\boldsymbol{V_1}(\boldsymbol{\beta}|\mathbf{X}) &= (\mathbf{V_0^{-1}} + \sigma^2\mathbf{Z'Z})^{-1}\\ &= \left[\begin{pmatrix}\mathbf{P_0^{-1}} & \mathbf{0}\\ \mathbf{0} & \mathbf{Q_0^{-1}}\end{pmatrix} + \begin{pmatrix}\mathbf{R^{-1}} & \mathbf{R^{-1}}\\ \mathbf{R^{-1}} & \mathbf{R^{-1}}\end{pmatrix}\right]^{-1},\end{aligned}$$

where $\mathbf{R} = \sigma^2\mathbf{(X'X)^{-1}}$ the least squares covariance matrix from a regression with $\mathbf{X}$. The posterior mean is,

$$E(\boldsymbol{\beta}|\mathbf{X}) = \mathbf{V_1(V_0}\boldsymbol{\beta_0} + \sigma^2\mathbf{Z'y).}$$

Our interest is in the posterior mean and variance for $\mathbf{b}$. Applying a standard formula for the inverse of a partitioned matrix (Harville 1997) to the posterior variance, the posterior covariance matrix for $\mathbf{b}$ is given by:

$$\boldsymbol{V_1}(\mathbf{b}|\mathbf{X}) = \mathbf{\Sigma + \Sigma R^{-1}[(R^{-1} + Q_0^{-1})^{-1} - R^{-1}\Sigma R^{-1}]R^{-1}\Sigma}, \qquad (A.4)$$

where $\mathbf{\Sigma} = \mathbf{(R^{-1} + P_0^{-1})^{-1}}$. We specify a diffuse prior for $\mathbf{b}$ in which $\mathbf{P_0^{-1} \approx 0}$ yielding a substantial simplification of equation (A.4):

$$\begin{aligned}\boldsymbol{V_1}(\mathbf{b}|\mathbf{X}) &= \mathbf{R + Q_0}\\ &= \sigma^2\mathbf{(X'X)^{-1} + Q_0}\end{aligned} \qquad (A.5)$$

With an improper diffuse prior for $\mathbf{b}$, the posterior variance is simply the least squares variance plus the prior covariance matrix for $\mathbf{a}$, the the potential bias in the OLS estimates. Because $\mathbf{R}$ and $\mathbf{Q_0}$ are positive definite, the posterior standard errors for $\mathbf{b}$ are larger than the OLS standard errors.

The posterior mean for $\mathbf{b}$ simplifies to

$$E(\mathbf{b}|\mathbf{X}) = \mathbf{(X'X)^{-1}X'y - a_0}. \qquad (A.6)$$

Not suprisingly, the posterior mean for $\mathbf{b}$ is simply the OLS estimate minus the prior of the bias in the OLS estimate.

If it were straightforward to think about the possible size and variance of $\mathbf{a}$ or equivalently the bias of OLS, we could simply stop here. In general, however, this is a difficult task. In the body of the paper we discuss ways of specifying priors for $\mathbf{c}$, the vector of covariances between $\mathbf{X}$ and $\mathbf{e}$. Going from a specification for a prior mean and covariance matrix for $\mathbf{c}$ to $\mathbf{a}$, however, is quite simple. Specifically:

$$\mathbf{a_0} = n\mathbf{(X'X)^{-1}c_0} \qquad (A.7)$$

2

and

$$\begin{aligned} V(\mathbf{a_0}) &= \mathbf{Q_0} \\ &= n^2(\mathbf{X'X})^{-1}\mathbf{U_0}(\mathbf{X'X})^{-1} \end{aligned} \qquad (A.8)$$

Substituting (A.8) into (A.5) yields:

$$\boldsymbol{V_1}(\mathbf{b}|\mathbf{X}) = \sigma^2(\mathbf{X'X})^{-1} + n^2(\mathbf{X'X})^{-1}\mathbf{U_0}(\mathbf{X'X})^{-1}$$

and substituting (A.7) into (A.6) yields:

$$E(\mathbf{b}|\mathbf{X}) = (\mathbf{X'X})^{-1}\mathbf{X'y} - n(\mathbf{X'X})^{-1}\mathbf{c_0}$$

We see here that the posterior mean for **b** is simply the OLS estimate of **b** adjusted by the prior value of the bias. The posterior covariance matrix for **b** is simply the OLS covariance matrix plus the covariance matrix associated with the bias.

3

## Appendix B. Computer Code

The following R code yields the Bayes estimates reported in the paper.

```
bayes <- function(x, y, u = 0.0001) {
# This function fits a Bayesian regression with a prior for model
# uncertainty. The prior parameter, u, describes the variance of a 0
# prior mean for the correlation between predictors and the error term.
# At the default value for u, correlations between X and e fall in the
# interval [-.02, .02].
#
# The user must supply a matrix of independent variables (without a
# constant) and a vector for the dependent variable. The function
# returns a matrix with 3 columns: (1) the OLS slope coefficients, (2)
# OLS standard errors, and (3) Bayesian standard errors.
#
   x <- scale(x, center = T, scale = F)  # Centering X
   y <- y - mean(y)                      # Centering y
   n <- length(y)                        # No. of obs.
   k <- ncol(x)                          # No. of X's (w/o intercept)
   U <- diag(rep(u, k))                  # Prior U matrix
   sdx <- sqrt(apply(x, 2, var))         # SD's of X
   x1 <- sweep(x, 2, sdx, "/")           # Scaling X
   ls <- summary(lm(y ~ x1 - 1))         # LS fit with scaled X
   sde <- sqrt(var(ls$res))              # SD of resids
   y1 <- y/sde                           # Scaling y
   ls <- summary(lm(y1 ~ x1 - 1))        # LS fit with scaled X and y
   xxi <- ls$cov                         # (X'X)^-1
   Q <- xxi %*% U %*% xxi * (n^2)        # Bayes variance adjustment
   lsv <- xxi * (ls$sigma^2)             # LS cov. matrix
   lsse <- sqrt(diag(lsv))               # LS SE's
   bv <- lsv + Q                         # Bayes cov. matrix
   bse <- sqrt(diag(bv))                 # Bayes SE's
   s <- sde/sdx                          # Rescaling vector
   b <- lm(y ~ x - 1)$coef               # unscaled LS coefs (w/o intercept)
   Table <- cbind(b, s * lsse, s * bse)  # Output (unscaled coefs and SE's)
   Table
}
```

4