Journal of Public Economics 128 (2015) 30–33

Contents lists available at ScienceDirect

Journal of Public Economics

journal homepage: www.elsevier.com/locate/jpube

War and local collective action in Sierra Leone: A comment on the use of coefficient stability approaches☆

Felipe González [a], Edward Miguel [a,b,*]

[a] Department of Economics, University of California, Berkeley, United States
[b] National Bureau of Economic Research

ARTICLE INFO

Article history:
Received 29 August 2014
Received in revised form 6 May 2015
Accepted 14 May 2015
Available online xxxx

JEL classification:
C4
D74

Keywords:
Selection
Unobservables
Civil war
Africa

ABSTRACT

In a study of the effect of civil war exposure on local collective action outcomes in Sierra Leone, Bellows and Miguel (2009) employ a coefficient stability approach to assess the importance of omitted variable bias building on Altonji et al. (2005a). Here we clarify the econometric assumptions underlying Bellows and Miguel (2009), and extend their analysis using data on dependent variable reliability ratios and the method developed in Oster (2015).

1. Introduction

Bellows and Miguel (2009) employ the Altonji et al. (2005a) coefficient stability approach to assess the importance of omitted variable bias, in a study of the effect of civil war exposure (called c below) on local collective action outcomes (y) in Sierra Leone. Specifically, they show that coefficient estimates on civil war exposure change little across regression specifications with and without additional covariates and, following a long tradition in applied economics, argue that omitted variable bias is unlikely to be driving their results.[1]

There is an implicit statistical assumption made in Bellows and Miguel (2009), namely, that the variances of the observable covariates ($x'\beta$) predicting local collective outcomes and of the unobservables ($\tilde{q}$) are equal:

$$\sigma_{x'\beta} = \sigma_{\tilde{q}}. \tag{1}$$

This is not exactly the same assumption that Altonji et al. (2005a) make in their analysis.[2] If Eq. (1) does not hold, then it follows that the stability of coefficient estimates across specifications with and without observed covariates need not represent convincing evidence of limited omitted variable bias.

As noted by Altonji et al. (2005a), the inherent difficulty that researchers face in establishing a plausible range of values for the ratio of $\sigma_{x'\beta}$ to $\sigma_{\tilde{q}}$ limits the practical utility of their approach. In a recent paper, Oster (2015) shows that if observables and unobservables have the same explanatory power in y (after taking into consideration any measurement error in y), then the following is a consistent estimator of the effect of c on y:

$$\hat{\hat{\alpha}} = \hat{\alpha}^* - (\hat{\alpha} - \hat{\alpha}^*) \times \frac{R_{max} - R^*}{R^* - R} \tag{2}$$

where $\hat{\alpha}^*$ and R^* are the coefficient estimate and R^2 from the regression including observable covariates, and $\hat{\alpha}$ and R are the coefficient and R^2 from the uncontrolled regression. In addition, R_{max} is the R^2 in a regression of y on all observable and unobservable controls, which is clearly unknowable (given its reliance on unobservables).

☆ We would like to thank Raj Chetty, Emily Oster, and three anonymous referees for comments and suggestions.

* Corresponding author at: Evans Hall #3880, University of California, Berkeley, CA 94720-3880
E-mail address: emiguel@berkeley.edu (E. Miguel).

[1] Beyond applying the approach in Altonji et al. (2005a), Bellows and Miguel also present a variety of other evidence that the estimated effects of civil war exposure on local collective action outcomes are causal.

[2] See condition 4 on page 175 of Altonji et al. (2005a), which states that "the relationship between [the treatment variable and the unobservables] is the same as the relationship between [the treatment variable and the observables], *after adjusting for differences in the variance of these distributions*" (our italics).

http://dx.doi.org/10.1016/j.jpubeco.2015.05.004

Although several recent papers have employed the approach laid out in Bellows and Miguel (2009), there is no compelling reason to assume that their assumption in Eq. (1) holds in general.[3] Therefore, the formula in Bellows and Miguel (2009) should not generally be applied to directly assess the degree of selection; instead, one should apply the corrected formula in Eq. (2) or the equivalent formulas in Oster (2015). Nevertheless, in order to apply this correction we need to have some sense of R_{max}, an unknown parameter.

The best we can hope for is to place plausible bounds on R_{max}. It is immediate that R^* is a lower bound on R_{max}, i.e., $R_{max} \in [R^*, 1]$. In addition, we note that R_{max} is bounded below one when there is classical measurement error in the dependent variable. Evidence on the extent of measurement error in y is thus a potentially useful way to generate an upper bound on R_{max}.[4] For example, in many low income country households datasets, it is well-known that income, consumption and business profits are measured with considerable error, in which case assuming that $R_{max} = 1$ or that it is close to 1 is likely to be far too conservative. For instance, McKenzie (2012) "demonstrates that for many economic outcomes, the autocorrelations are typically lower than 0.5, with many around 0.3", with values ranging from 0.2 to 0.8. McKenzie (2012) and De Mel et al. (2009) both suggest that measurement error in the outcome variables is substantial. As a point of reference, in relatively high quality U.S. survey data Angrist and Krueger (1999) conclude that reliability ratios are typically between 0.7–0.9 for the most commonly studied labor market outcomes.

One approach to begin quantifying the extent of measurement error is through the reliability ratio for a variable that should be fixed over time. For example, if the survey–resurvey reliability ratio of an outcome variable is 0.8 (such as in reliably measured labor market outcomes in U.S. datasets), then the maximum attainable R^2 (with any set of controls) is just 0.8.[5] In what follows we apply the procedure in Oster (2015), using survey–resurvey reliability ratio data, to adjust the coefficient estimates in Tables 3, 4, and 5 in Bellows and Miguel (2009).

2. Impacts of war exposure

2.1. Bounds on R_{max}

The formula in Eq. (2) delivers much wider bounds on the potential bias due to unobservables. In order to obtain more informative bounds, one can gauge the amount of residual variance that comes purely from measurement error by using multiple measures of the same quantity. We adjust coefficients in Bellows and Miguel (2009) using four different approaches: (1) the Bellows and Miguel approach ($R_{max} = 2R^*–R$), (2) the reliability ratio approach developed in this paper, (3) the Oster approach ($R_{max} = \min\{2.2R^*, 1\}$), and (4) the most conservative case ($R_{max} = 1$). Each of these approaches gives different estimates of R_{max} and, therefore, a different estimate of the bounds.

As we need estimates of the reliability ratio in the second approach, it deserves additional explanation. Although there is no information on reliability for the Sierra Leone data used in Bellows and Miguel (2009), Baird et al. (2008) contains an analysis of survey–resurvey data from Kenya.[6] By focusing on variables where we have reason to believe that responses should be unchanging over time, the setup in Baird et al. (2008) allows them to estimate how much "noise" there is in reported outcomes in survey data in a rural African sample. We focus on the educational attainment of the respondent's father as a variable that should be fixed over time and where any variation should presumably be due to reporting error.[7] As reported in Baird et al. (2008), the correlation between stated father's educational attainment across the two survey rounds collected three months apart was moderate: the pairwise correlation coefficient is 0.80. This implies that if father's educational attainment were an outcome variable, the maximum attainable R^2 (with any set of controls) would be just 0.8. Therefore, we use $R_{max} = 0.8$ when obtaining bounds using the reliability ratio approach. We also present the results using the reliability ratio approach using the less conservative assumption of 0.5, roughly based on the data in McKenzie (2012).

2.2. Results

Table 1 first reproduces results for the key coefficient on civil war exposure in Bellows and Miguel (2009) in columns 1 and 2. We do this for each of the three outcome variables presented in Tables 3, 4 and 5 of Bellows and Miguel, namely, community meeting attendance in the last year (the top panel of Table 1); membership in a social group (the middle panel); and membership in a political group (the bottom panel). We then proceed to present five versions of the results in Bellows and Miguel (2009) using different statistical assumptions regarding R_{max}, based on the discussion above. Building on Oster (2015), the result of each extension is an interval of values for the war exposure coefficient that are consistent with the degree of omitted variable bias accommodated by the R_{max} value, the difference between the R^2 in the uncontrolled and controlled regression, and the change in the estimated coefficient across those specifications.

Specifically, columns 1 and 2 present results from the uncontrolled and the controlled regressions, respectively, in Bellows and Miguel (2009), and the corresponding R^2 values. The increase in the R^2 across these two columns captures the amount of variation in the dependent variable that is explained by the observed covariates. This, together with the maximum amount of variation that can be potentially explained (R_{max}), gives rise to an estimate of the lower bound for the interval of coefficients. If this lower bound is greater than zero, it would provide further evidence that the true causal effect of civil war exposure on the outcome is indeed likely to be positive.

In column 3, we again reproduce the analysis in Bellows and Miguel (2009) using the R_{max} implied by their assumption in Eq. (1) laid out above. In this case, the R_{max} is simply the R^2 from the controlled regression plus the difference between the R^2 in the uncontrolled and controlled regressions, i.e., $R_{max} = R^* + (R^*–R)$. The intervals of coefficient values in this case are equivalent to the results presented in the original paper for all three outcome variables, with intervals containing strictly positive values for all three outcome variables.[8]

The next set of results in columns 4 and 5 generate different values of R_{max} using the reliability ratio figure in Baird et al. (2008), namely, $R_{max} = 0.8$ in column 4, and the assumption of $R_{max} = 0.5$ in McKenzie (2012) (both are higher values of R_{max} than that implicitly assumed in BM-09). The interval of coefficient values is broader as a result, increasing from [0.060, 0.065] to [0.048, 0.065] for the community meeting attendance indicator, although even in this case, the lower bound remains greater than zero, providing some evidence of a positive

[3] See Cavalcanti et al. (2010), Essaji and Fujiwara (2012), de Brauw and Mueller (2012), Hermes et al. (2012), Sampaio et al. (2013), Breuer and McDermott (2013), Minoiu and Shemyakina (2014), Wong (2014), Jiraporn et al. (2014), and Kosec (2014), among others.

[4] Oster acknowledges the practical difficulty in establishing an empirically grounded value for R_{max}. The use of prior information about the determinants of observed y to inform the extent of bias has been previously suggested by Altonji et al. (2005a,b, 2008).

[5] To see this more clearly, note that in the presence of measurement error ε the maximum R^2 in a regression of y on observables x and unobservables $\tilde{q}$ is $R_{max} = \mathrm{Var}(\alpha c + x'\beta + \tilde{q})/\mathrm{Var}(\alpha c + x'\beta + \tilde{q} + \varepsilon)$, which corresponds to the reliability ratio.

[6] We feel the use of estimates from another rural African setting is a useful starting point for illustrating the approach, and can produce suggestive results, although it would clearly be preferable to use data from within the study sample itself.

[7] In our view, this is an appropriate variable to establish a plausible upper bound on the reliability ratio because it should be stable over time, relatively simple, and is salient in the local context. As noted above, this is just suggestive since it comes from a different setting than the data used in Bellows and Miguel (2009).

[8] Bellows and Miguel (2009) employ a moderate number of covariates, as described in the note to Table 1.

Table 1
Impacts of war exposure.

	(1)	(2)	(3)	(4)	(5)	(6)	(7)
			Bellows and Miguel (2009)	Reliability ratio		Oster (2015) approach	Most conservative case
				McKenzie (2012)	Baird et al. (2008)		
Table 3							
Did you attend any community meetings in the past year?							
Conflict victimization index	0.070***	0.065***	[0.060, 0.065]	[0.048, 0.065]	[−0.001, 0.065]	[−0.011, 0.065]	[−0.033. 0.065]
	(0.018)	(0.017)					
R^2	0.362	0.391	–	–	–	–	–
R_{max}	–	–	0.420	0.500	0.800	0.860	1.000
Table 4							
Are you a member of a social group?							
Conflict victimization index	0.071***	0.066***	[0.059, 0.066]	[0.012, 0.066]	[−0.092, 0.066]	[−0.077, 0.066]	[−0.161, 0.066]
	(0.018)	(0.018)					
R^2	0.328	0.345	–	–	–	–	–
R_{max}	–	–	0.362	0.500	0.800	0.759	1.000
Table 5							
Are you a member of a political group?							
Conflict victimization index	0.061***	0.057***	[0.053, 0.057]	[−0.008, 0.057]	[−0.100, 0.057]	[−0.050, 0.057]	[−0.162. 0.057]
	(0.013)	(0.013)					
R^2	0.276	0.289	–	–	–	–	–
R_{max}	–	–	0.302	0.500	0.800	0.636	1.000
Controls	No	Yes					
Area/Year F.E.	Yes	Yes					
Observations	10,471	10,471					

Notes. Dependent variable in *italics*. Set intervals estimation using an equal proportional selection assumption. Oster's approach uses $R_{max} = 2.2 \times R^*$, where R^* is the R^2 from column (2). Controls include indicator variables for females, "Respondent has any education", "Traditional authority household", and an age variable.

*** $p < 0.01$.
** $p < 0.05$.
* $p < 0.1$.

effect. For the social group membership and political group membership variables, the range of the interval now is close to zero in column 4 and includes zero in column 5. Thus adopting a more conservative approach to R_{max} than that implicitly used in Bellows and Miguel (2009) we cannot rule out the hypothesis that omitted variable bias drives their results in the case of $R_{max} = 0.8$, while under the less conservative assumption of $R_{max} = 0.5$ the bounds do not generally include zero.

We next present two further bounding approaches. Column 6 is based on the parameterization suggested by Oster (2015) in her empirical cases, namely, $R_{max} = \min\{\Pi R^*, 1\}$, where Π is empirically estimated to equal roughly 2.2; Oster herself is explicit about the fact that this parameter estimate for Π may not apply to other settings, but the goal of this column is to provide an illustration of her approach. The most conservative possible approach is presented in column 7, under the assumption that $R_{max} = 1$. This latter assumption implicitly assumes that there is exactly zero measurement error in the reported outcome variable, i.e., perfect survey–resurvey reliability, which seems implausible in most real-world applications, especially with low-income country survey data. Under both of these sets of assumptions, the interval of coefficient values increases, and in the latter case is so large as to become largely uninformative. Fig. 1 presents these results graphically.

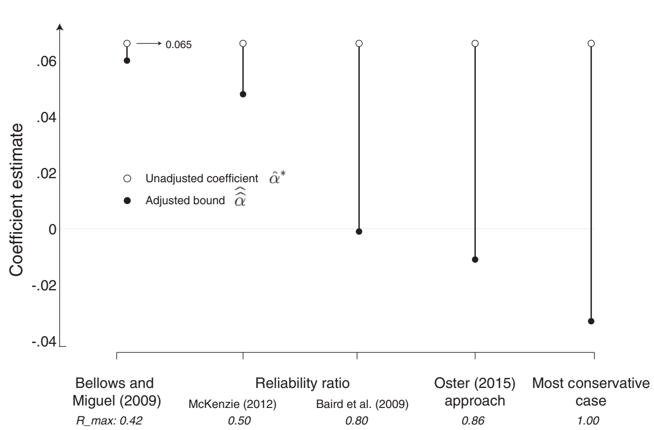

Fig. 1. Impacts of war exposure. *Notes*: Impacts of war exposure under different assumptions about R_{max}, i.e., the maximum amount of variation that can be explained in a regression of a dependent variable of interest on war exposure. The dependent variable is an indicator variable that takes the value of one if the answer to the following question is yes: "Did you attend any community meetings in the past year?".

It is apparent that the reliability ratio method to determining a plausible R_{max} yields a more conservative bound than that generated by the approach in Bellows and Miguel (2009), but one that may be more informative than the approach in Oster (2015) or than the unrealistically conservative assumption of $R_{max} = 1$. An R_{max} value based on the precise variables used in Bellows and Miguel's Sierra Leone study (rather than the values from other low income settings that we use out of convenience here) would clearly be necessary to place more definitive bounds on the coefficient estimates in Bellows and Miguel (2009), but the illustrative results in Table 1 of this note do, we hope, make the case that the proposed reliability ratio approach is a promising way forward.

3. Discussion

This note illustrates a method to operationalizing recent coefficient stability approaches to assessing the degree of omitted variable bias, namely those in Altonji et al. (2005a) and Oster (2015), with an application to the analysis in Bellows and Miguel (2009). The approaches in all three of these papers rely on untested assumptions about the nature of unobserved explanatory variables, which considerably limits their applicability. We have illustrated several ways to apply these tools using different assumptions about R_{max}.

However, the data needed to put bounds on R_{max} is rarely found in most current datasets. Yet as shown in Baird et al. (2008) and McKenzie (2012), this data is not very difficult to collect in most field survey data collection exercises, like those that are now common in development economics and increasingly common in other economics sub-fields. The random sampling of a subset of respondents for resurveys, who would then shortly afterwards be re-asked the questions needed to generate the study's leading outcome variables, would typically be sufficient to estimate reliability and to quantitatively assess the amount of “noise” in key variables of interest. Since it is likely that R_{max} differs considerably across empirical settings, datasets, and particular variables, this information would need to be routinely collected for the variables of interest in order to be most credible (although a range of typical reliability ratio values would likely emerge over time if enough such survey–resurvey data is collected, at least for commonly studied variables). Without a better understanding of what a plausible upper bound on R_{max} might be for a given dependent variable, the approaches in Altonji et al. (2005a) and Oster (2015) will often be of limited empirical applicability. Armed with this information, however, these econometric methods could become more powerful tools for addressing omitted variable bias concerns.

References

Altonji, Joseph G., Elder, Todd E., Taber, Christopher R., 2005a. Selection on observed and unobserved variables: assessing the effectiveness of catholic schools. J. Polit. Econ. 113 (1), 151–184.

Altonji, Joseph G., Elder, Todd E., Taber, Christopher R., 2005b. An evaluation of instrumental variable strategies for estimating the effects of catholic schooling. J. Hum. Resour. XL (4), 791–821.

Altonji, Joseph G., Elder, Todd E., Taber, Christopher R., 2008. Using selection on observed variables to assess bias from unobservables when evaluating Swan-Ganz catheterization. Am. Econ. Rev. Pap. Proc. 98 (2), 345–350.

Angrist, Joshua, Krueger, Alan, 1999. Empirical strategies in labor economics. In: Ashenfelter, Orley, Card, David (Eds.), The Handbook of Labor Economics vol. III. North Holland (chapter 23).

Baird, Sarah, Hamory Hicks, Joan, Miguel, Edward, 2008. Tracking, Attrition and Data Quality in the Kenyan Life Panel Survey Round 1 (KLPS-1). University of California CIDER Working Paper.

Bellows, John, Miguel, Edward, 2009. War and local collective action in Sierra Leone. J. Public Econ. 93, 1144–1157.

Breuer, Janice Boucher, McDermott, John, 2013. Respect, responsibility, and development. J. Dev. Econ. 105, 36–47.

Cavalcanti, Tiago, Guimaraes, Juliana, Sampaio, Breno, 2010. Barriers to skill acquisition in Brazil: public and private school students performance in a public university entrance exam. Q. Rev. Econ. Financ. 50, 395–407.

de Brauw, Alan, Mueller, Valerie, 2012. Do limitations in land rights transferability influence mobility rates in Ethiopia. J. Afr. Econ. 21 (4), 548–579.

De Mel, Suresh, McKenzie, David J., Woodruff, Christopher, 2009. Measuring microenterprise profits: must we ask how the sausage is made? J. Dev. Econ. 88, 19–31.

Essaji, Azim, Fujiwara, Kinya, 2012. Contracting institutions and product quality. J. Comp. Econ. 40, 269–278.

Hermes, Niels, Kihanga, Ernest, Lensink, Robert, Lutz, Clemens, 2012. The impact of trade credit on customer switching behaviour: evidence from the Tanzanian rice market. J. Dev. Stud. 48 (3), 363–376.

Jiraporn, Pornsit, Liu, Yixin, Kim, Young S., 2014. How do powerful CEOs affect analyst coverage? Eur. Financ. Manag. 20 (3), 652–676.

Kosec, Katrina, 2014. The child health implications of privatizing Africa's urban water supply. J. Health Econ. 35, 1–19.

McKenzie, David, 2012. Beyond baseline and follow-up: the case for more t in experiments. J. Dev. Econ. 99 (2), 210–221.

Minoiu, Camelia, Shemyakina, Olga N., 2014. Armed conflict, household victimization, and child health in Cote d'Ivoire. J. Dev. Econ. 108, 237–255.

Oster, Emily, 2015. Unobservable Selection and Coefficient Stability: Theory and Evidence.

Sampaio, Breno, Sampaio, Gustavo Ramos, Sampaio, Yony, 2013. On estimating the effects of immigrant legalization: do U.S. agricultural workers really benefit. Am. J. Agric. Econ. 1–17.

Wong, Maisy, 2014. Estimating the distortionary effects on ethnic quotas in Singapore using housing transactions. J. Public Econ. 115, 131–145.