



University of
Massachusetts
Amherst

Bounding Sets for Treatment Effects with Proportional Selection

Item Type	Working Paper
Authors	Basu, Deepankar
DOI	10.7275/22680948
Rights	UMass Amherst Open Access Policy
Download date	2025-06-13 18:35:54
Link to Item	https://hdl.handle.net/20.500.14394/22292

Bounding Sets for Treatment Effects with Proportional Selection

Deepankar Basu*

April 15, 2021

Abstract

In linear econometric models with proportional selection on unobservables, omitted variable bias in estimated treatment effects are roots of a cubic equation involving estimated parameters from a short and intermediate regression, the former excluding and the latter including all observable controls. The roots of the cubic are functions of δ , the degree of proportional selection on unobservables, and R_{max} , the R-squared in a hypothetical long regression that includes the unobservable confounder and all observable controls. In this paper a simple method is proposed to compute roots of the cubic over meaningful regions of the δ - R_{max} plane and use the roots to construct bounding sets for the true treatment effect. The proposed method is illustrated with both a simulated and an observational data set.

Keywords: treatment effect, omitted variable bias.

JEL Codes: C21.

1 Introduction

Researchers are often interested in estimating treatment effects in models where there are clear problems of unobserved or unobservable confounders. To fix ideas, consider the following linear regression model,

$$Y = \beta X + \Psi \omega^0 + W_2 + \varepsilon, \quad (1)$$

where Y is the scalar outcome variable, X is the scalar treatment variable of interest, ω^0 is a $J \times 1$ vector of observed controls, Ψ is $1 \times J$ vector of

*Department of Economics, University of Massachusetts Amherst. Email: dbasu@econs.umass.edu. I would like to thank Leila Gautham for comments on an earlier version of this paper.

parameters, and W_2 is an unobserved confounder. Suppose a researcher is interested in consistently estimating β , but is unable to do so because of the presence of the unobservable confounder, W_2 (which can be thought of as an index of a set of unobservable variables), in the hypothetical ‘long’ regression model.

Faced with this problem, researchers often compare the ordinary least square (OLS) estimate of β between a ‘short’ and an ‘intermediate’ regression, where the short regression is given by

$$Y = \hat{\beta}X + \hat{\varepsilon}, \tag{2}$$

in which both the observable and unobservable controls, i.e. ω^0 and W_2 , are missing from the model, and the intermediate regression is given by

$$Y = \tilde{\beta}X + \tilde{\Psi}\omega^0 + \tilde{\varepsilon}, \tag{3}$$

in which only the unobservable control, W_2 , is missing from the model. If the numerical magnitude of $\hat{\beta}$ and $\tilde{\beta}$ are roughly similar, i.e. the estimate of the treatment effect is ‘stable’, researchers conclude that the bias from the omitted, unobservable confounder is small.

In a recent, innovative contribution, [Oster \(2019\)](#) has demonstrated that such ‘coefficient stability’ arguments to deal with possible omitted variable bias is misleading.¹ In fact, what is needed to draw conclusions about the magnitude of possible bias due to the unobservable confounder is not the raw change in the estimate of the treatment effect, but an R-squared scaled change in the estimate of the treatment effect between the short and intermediate regressions. This becomes clear when we write the expression for the omitted variable bias in the OLS estimate of the treatment effect in the intermediate regression in terms of the R-squared in the short, intermediate and long regressions, and relevant coefficients in the long regression. A little algebraic manipulation generates a cubic equation in the bias (of the OLS estimate of the treatment effect in the intermediate regression).

A cubic equation with real coefficients will have either one or three real roots. When the cubic equation has a unique real root, the researcher is able to identify the bias, and hence the bias-adjusted treatment effect, without any ambiguity. When the cubic equation has three real roots, the researcher is confronted with the problem of non-uniqueness. [Oster \(2019\)](#) proposes two approaches to deal with the problem of non-uniqueness.

The first method involves computing the bias-adjusted treatment effect under the twin assumptions of $\delta = 1$ (equal selection on observables and

¹[Oster \(2019\)](#) extends previous work on this issue by [Altonji et al. \(2000, 2005\)](#).

unobservables) and a sign restriction (which is stated as Assumption 3 in her paper). In this case, [Oster \(2019, pp. 194\)](#) argues, we can arrive at a unique solution for the bias in the treatment effect and can therefore compute a unique bias-adjusted treatment effect.

The second method relies on choosing some value of R_{max} (the magnitude of R-squared in a hypothetical long regression that includes all variables, including the unobservable confounder), and calculating the magnitude of δ , i.e. degree of selection due to unobservables, that would be consistent with $\beta = 0$ (no treatment effect). In this case, [Oster \(2019\)](#) shows that we are able to find a unique magnitude of proportional selection that would make the treatment effect vanish.

Both these methods promise to be enormously useful for applied researchers because they provide workable solutions for the pervasive and rather intractable problem of omitted variable bias ([Basu, 2020](#)). That is why the method proposed by [Oster \(2019\)](#) has been widely cited in economics and the social sciences.² Unfortunately, as I demonstrate in this paper, both methods suffer from serious problems.

The first method, i.e. computing bias-adjusted treatment effect under the assumption of equal selection, suffers from many theoretical problems. First, without additional assumptions, it is not possible to ensure the existence of a unique solution. But these assumptions cannot be justified either on theoretical or empirical grounds. Second, once these assumptions are imposed, there is no leeway for researchers to experiment with different values of R_{max} because a specific value of R_{max} gets pinned down. Third, the method will not work for cases where estimates of the treatment effect declines with the addition of control variables. Finally, there is a sharp discontinuity at $\delta = 1$, i.e. conclusions can change dramatically if δ is perturbed even slightly from the value of unity.

The second proposal in [Oster \(2019\)](#) is to compute the value of δ that is consistent with a zero treatment effect. This method suffers from two problems. First, the method relies on choosing a unique value of R_{max} . Since there is no reliable way to pin down a unique R_{max} (the R-squared in the hypothetical long regression), it reduces robustness of the method. Second, while this method allows us to compute a unique value of $\delta = \delta^*$ (proportion of selection due to unobservables) that is associated with a zero treatment effect, there seems to be some misunderstanding about how to interpret this

²On Professor Oster’s google scholar page, the paper shows 1611 citations. Here are just a few examples: [Galor and Ozak \(2016\)](#); [Michalopoulos and Papaioannou \(2016\)](#); [Goldsmith-Pinkham et al. \(2019\)](#); [Jaschke and Keita \(2021\)](#). Papers published before 2019 cite different working paper versions of [Oster \(2019\)](#).

δ^* . Altonji et al. (2005) and Oster (2019) use informal arguments to argue that $\delta = 1$ should be used as a lower bound. Hence, this method recommends that if the computed value of δ^* is higher than unity, then researchers can conclude that their reported results do not suffer from omitted variable bias (the larger the δ^* the better); if the computed value of δ^* is lower than unity, then researchers should be worried about omitted variable bias (the smaller the δ^* the more serious the problem).

A computed value of $\delta^* > 1$ means that even if the unobservables are more strongly correlated to the treatment variable than the observables, the omitted variable bias cannot wipe out the reported nonzero treatment effect. Thus, this methodology suggests, a computed value of $\delta > 1$ provides confidence in the results. This conclusion does not seem to be warranted. A value of $\delta^* > 1$ does not tell us anything about the magnitude of bias if $\delta < 1$. The implicit understanding seems to be that any $\delta < \delta^*$ would not be inimical to the reported result. But this is not true. Even if $\delta^* > 1$, it is possible for omitted variable bias to completely nullify the treatment effect for a value of $\delta < 1$ (I show this in an actual example in section 5.2). In a similar manner a computed value of $\delta^* < 1$ does not provide evidence that the reported result will necessarily be nullified if we take account of omitted variable bias (I show this in an actual example in section 5.1). Thus, while the computation of δ^* is straightforward, it is not very informative, which goes completely against the grain of the existing literature that has used δ^* widely.

I propose an alternative method to to quantify the bias in the treatment effect that significantly improves on Oster’s methodology. To understand my proposal let us return to the cubic equation that is at the heart of the bias calculations. The coefficients of the cubic equation are functions of two unknown parameters: R_{max} , the R-squared in the hypothetical long regression that includes all observable and unobservable variables, and δ , the relative degree of selection on unobservables.

The first step in my proposal is to choose a meaningful bounded box in the δ - R_{max} plane and divide it into two parts, the first corresponding to unique real roots of the cubic equation and the second corresponding to nonunique real roots. Let us call the first the *URR* (unique real root) area and the second the *NURR* (nonunique real root) area. Both regions are function of δ and R_{max} , i.e. $URR = URR(\delta, R_{max})$, and $NURR = NURR(\delta, R_{max})$. On each point in $URR(\delta, R_{max})$ one can compute a unique real root of the cubic equation. I do so on a sufficiently granular $N \times N$ grid of this area, *URR*, and collect the real roots in a vector, B_U , (whose length is equal to N^2 , the number of points of the grid).

Proposition 2 in [Oster \(2019\)](#) shows that the ‘true’ treatment effect is the difference in the treatment effect estimated from the intermediate regression, $\tilde{\beta}$, and the root of the cubic, ν . Thus, the difference between $\tilde{\beta}$ and each element of the vector, B_U , gives me a consistent estimate of the ‘true’ treatment effect.

In the second step, I compute a bounding set for the ‘true’ treatment effect as the interval formed by the 2.5-th and the 97.5-th quantile of the empirical distribution of $\beta^* = \tilde{\beta} - \nu$. This bounding set contains the ‘true’ treatment effect with 95% probability. If this bounding set does not include zero, a researcher can conclude that the reported treatment effect will not be wiped out even after we have taken account of possible omitted variable bias. If this bounding set contains zero, then that raises concerns about the reported treatment effect.

In the final step, I turn to the *NURR* area. On each point in the *NURR* area, the cubic equation has three real roots. I compute the three roots of the cubic on a sufficiently granular grid of *NURR*. Since Proposition 2 in [Oster \(2019\)](#) shows that one of these three real roots is the true bias in the treatment effect, I need to choose one of them. Assuming that what is known - the distribution of the root on the *URR* area being the true bias - can provide information about what is not known - the distribution of the correct root on the *NURR* area, I choose the root whose distribution is ‘closest’ to the distribution of the root on the *URR* area - ‘close’ in the sense of location of the two distribution. There can be many ways to implement the notion of closeness of the distributions. I choose to use a measure of central tendency - the median - to choose among the three roots computed on the *NURR* area, i.e. I choose the root whose empirical distribution has a median that is closest to the median of the empirical distribution of the root computed on the *URR* area. I recompute the bounding set using the chosen root from the *NURR* area and see if my earlier results change. If they do not, then this provides evidence of robustness. If the results are drastically different, e.g. one bounding set contains zero while the other does not, then I conclude that the method proposed in this paper is not applicable for that particular analysis.

My proposed method has clear advantages over the methodology proposed in [Oster \(2019\)](#). In comparison to the first method of [Oster \(2019\)](#), my method is free of the theoretical problems that I identify in her method. My method offers a cleaner, theoretically grounded method of computing bounds for the ‘true’ treatment effect. In comparison to the second method proposed by [Oster \(2019\)](#), my method is more robust but also more conservative. Instead of choosing a specific value of R_{max} , I compute and then use

the treatment bias for all possible values of R_{max} and δ over a meaningful area.

After presenting the theoretical results, I use my method on two data sets to illustrate the methodology. Both these data sets have been used in Oster (2019) and using these data facilitates easy comparison of our respective methodologies.³ The first data set contains simulated data on earnings using the NLSY-79 cohort. I use it to run Mincerian wage regressions and estimate returns to education. I compare my method with the one used in Oster (2019, section 4.1). The second data set comes from the Children and Young Adults sample of the NLSY and is used to study the impact of maternal behaviour on child outcomes. Using this data set, I highlight the differences in my methodology from the results reported and discussed in Oster (2019, section 4.2).

The rest of the paper is organized as follows. In section 2, I discuss the basic set-up; in section 3, I present my method of analyzing bias; in section 4, I discuss the disadvantages of working with the assumption of equal selection; in section 5, I illustrate my method, and contrast my results with Oster’s method, using two data sets, a simulated data set (NLSY data set to investigate the returns to education) and an actual data set (NLSY data set to study the impact of maternal behaviour on child outcomes); in section 6, I conclude with a summary of my proposed methodology for applied researchers and highlight a weakness of my proposed methodology.

2 Basic Set-Up

2.1 Three Regression Models

Consider once again the hypothetical ‘long’ regression,

$$Y = \beta X + \Psi\omega^0 + W_2 + \varepsilon, \quad (4)$$

and denote by R_{max} , the R-squared from the long regression. Consider the ‘short’ regression,

$$Y = \hat{\beta}X + \hat{\varepsilon}, \quad (5)$$

and denote as \hat{R} , the R-squared from the short regression. In a similar manner, consider an intermediate regression,

$$Y = \tilde{\beta}X + \tilde{\Psi}\omega^0 + \tilde{\varepsilon}, \quad (6)$$

³I would like to thank Emily Oster for making her data set available. I have downloaded the data sets from her webpage: <https://drive.google.com/file/d/OB1U4uS7GkkxbV0VkZmd0ZV1DVDA/view?usp=sharing>

and denote by \tilde{R} , the R-squared from the intermediate regression. Note that $R_{max} \geq \tilde{R} \geq \hat{R}$ (Greene, 2012, Theorem 3.6, pp. 42). Finally, consider an auxiliary regression,

$$X = \alpha\omega^0 + u, \quad (7)$$

and denote by \tilde{X} , the residual from this auxiliary regression. Let $\hat{\tau}_X$ denote the variance of \tilde{X} , σ_X^2 denote the variance of X and σ_Y^2 denote the variance of Y , and note that $\sigma_X^2 > \tau_X$.

2.2 Proportion of Selection

Following Oster (2019), let us define the measure of proportional selection on unobservables as,

$$\delta = \frac{\sigma_{2X}/\sigma_2^2}{\sigma_{1X}/\sigma_1^2} \quad (8)$$

where $\sigma_{1X} = cov(W_1, X)$, $\sigma_{2X} = cov(W_2, X)$, $\sigma_1^2 = var(W_1)$, and $\sigma_2^2 = var(W_2)$, and $W_1 = \Psi\omega^0$ (an index of the observable controls). Let us try to understand the meaning of this parameter, δ ?

Consider a linear projection (Wooldridge, 2002, chapter 2) of the treatment variables on the index of the observables, i.e.

$$X = \alpha_0 + \alpha_1 W_1 + u_1.$$

Since u_1 is orthogonal to W_1 by definition of linear projections, we have

$$\alpha_1 = \frac{\sigma_{1X}}{\sigma_1^2}. \quad (9)$$

Now consider another linear projection of the treatment variables on the index of the unobservables, i.e.

$$X = \delta_0 + \delta_1 W_2 + u_2$$

and note, once again using the property of linear projections, that

$$\delta_1 = \frac{\sigma_{2X}}{\sigma_2^2}. \quad (10)$$

Now we see clearly that the measure of proportional selection is just the ratio of the two coefficients from the two linear projections, i.e.

$$\delta = \frac{\delta_1}{\alpha_1}. \quad (11)$$

We will return to this expression when we try to look critically at the use of $\delta = 1$ as a lower bound.

2.3 Solving a Cubic

Let us denote by ν the bias in the treatment effect estimated from the intermediate regression. Using the well-known formula for omitted variable bias in the short and intermediate regressions, [Oster \(2019, Appendix A.4\)](#) shows that, in large samples,

$$\left(\mathring{\beta} - \tilde{\beta}\right) = \frac{\sigma_{1X}}{\sigma_2^2} - \nu \left(\frac{\sigma_X^2 - \tau_X}{\sigma_X^2}\right). \quad (12)$$

Recall that R-squared is the ratio of the squared fitted values and the square of the outcome variable. The fitted value is the product of the estimated regression coefficient and the regressors. The estimated regression coefficient in the short and intermediate regressions include the bias, ν . Hence, for the short and intermediate regressions, the expressions for the R-squared will have the treatment bias in the numerators. Using this idea, [Oster \(2019, Appendix A.4\)](#) shows that, in large samples, we will have

$$\left(\tilde{R} - \mathring{R}\right) = \sigma_1^2 + \tau_X \nu^2 - \frac{1}{\sigma_X^2} (\sigma_{1X} + \nu \tau_X)^2 \quad (13)$$

and

$$\left(R_{max} - \tilde{R}\right) = \nu \left(\frac{\sigma_1^2 \tau_X}{\sigma_{1X}} - \nu \tau_X\right). \quad (14)$$

The equations in (12), (13) and (14) constitute a system of 3 equations in 3 unknowns: σ_1^2 , the variance of W_1 ; σ_{1X} , the covariance of W_1 and X (treatment variable); and ν (the bias of the treatment effect in the intermediate regression). Algebraic manipulation can reduce the three equations into a single cubic equation in ν given by,

$$a\nu^3 + b\nu^2 + c\nu + d = 0, \quad (15)$$

where

$$a = (\delta - 1) (\tau_X \sigma_X^2 - \tau_X^2) \quad (16)$$

$$b = \tau_X \left(\mathring{\beta} - \tilde{\beta}\right) \sigma_X^2 (\delta - 2) \quad (17)$$

$$c = \delta \left(R_{max} - \tilde{R}\right) \sigma_Y^2 (\sigma_X^2 - \tau_X) - \left(\tilde{R} - \mathring{R}\right) \sigma_Y^2 \tau_X - \sigma_X^2 \tau_X \left(\mathring{\beta} - \tilde{\beta}\right)^2 \quad (18)$$

$$d = \delta \left(R_{max} - \tilde{R}\right) \sigma_Y^2 \left(\mathring{\beta} - \tilde{\beta}\right) \sigma_X^2 \quad (19)$$

3 Bounds for the Treatment Effect

3.1 Real Root as Bias

Finding the real roots of the cubic equation in (15) is the key to constructing proper bounds for the ‘true’ treatment effect. This follows from Proposition 2 in Oster (2019). In the case when there is only one real root, denote it by ν_1 . If $\beta^* = \tilde{\beta} - \nu_1$, then $\beta^* \xrightarrow{P} \beta$, so that ν_1 is the asymptotic bias in the treatment effect estimated by the intermediate regression. Hence, in large samples, $\tilde{\beta} - \nu_1$ is the bias-adjusted treatment effect. In the case when the cubic equation has three real roots, ν_1, ν_2, ν_3 , then only one of these will give us the asymptotic bias in the treatment effect. Hence, only one of the following, $\tilde{\beta} - \nu_1$, $\tilde{\beta} - \nu_2$ and $\tilde{\beta} - \nu_3$, will be the bias-adjusted treatment effect.

Our main task in constructing bounding sets for the true treatment effect is to identify the relevant real roots of cubic equation in (15). To do so we note that the coefficients of the cubic equation are composed of all known quantities other than the following two: R_{max} (the R-squared in the hypothetical long regression), and δ (the measure of proportional selection on unobservables). Therefore, our primary strategy will be to identify the area of the (δ, R_{max}) plane where the cubic (15) is guaranteed to have a unique real root, compute all real roots on that area, and use these roots to define β^* . As a secondary strategy, we will use one of the non-unique real roots to check for robustness of our results.

3.2 Unique Real Root and the Bounding Set

The nature of the roots of a cubic equation depend on the sign of its discriminant: if the discriminant is positive, there is a unique real root; if the discriminant is nonpositive, there are three real roots (for details, see the appendix). For the cubic equation in (15), let

$$p = \frac{3ac - b^2}{3a^2}, \tag{20}$$

$$q = \frac{27a^2d + 2b^3 - 9abc}{27a^3}; \tag{21}$$

then the discriminant of the cubic equation is given by $(27q^2 + 4p^3) / 108$.

Proposition 1. *If $27q^2 + 4p^3 > 0$, then the cubic equation in (15) will have a unique real root.*

Proof. This is a well known result. Please see the appendix for some discussion. \square

We will first demarcate a bounded box on the (δ, R_{max}) plane and then use the condition in Proposition 1 to identify the *URR* (unique real root) area in that box. Figure 1 provides a visual representation of the procedure, where we measure δ on the x-axis and R_{max} on the y-axis. We start by noting that R_{max} lies between \tilde{R} (because the hypothetical long regression has more regressors than the intermediate regression) and 1; hence, the bounded box must have $\tilde{R} \leq R_{max} \leq 1$. The next task is to choose a bounded interval on the δ axis, $\delta_{min} \leq \delta \leq \delta_{max}$. Once the researcher has chosen δ_{min} and δ_{max} , the bounded box, *ABCD* is in place.

To make the analysis robust, the researcher must choose δ_{min} to be a small positive number, and δ_{max} to be a large positive number. Note that $\delta = 0$ is not interesting because this is the assumption that the omitted variable (or index of omitted variables) is not correlated with the treatment). This seems to be unlikely and hence we must choose δ_{min} to be a small positive number, e.g. $\delta_{min} = 0.01$. On the other end, δ_{max} captures the maximum possible strength of relative selection on unobservables. For robustness, the researcher should choose a large positive number, e.g. $\delta_{max} = 5$. This would mean that the research will allow for the possibility that selection on unobservables is up to five times stronger than selection on observables.

[Figure 1 about here]

Once the bounded box, *ABCD*, is in place, we will trifurcate it. We do so for two reasons. First, we need to keep the important point in mind that at $\delta = 1$, the cubic equation in (15) is no longer a cubic (because the coefficient on ν^3 becomes zero). Hence, the researcher needs to bound the δ -intervals away from 1. To highlight this, I have indicated the vertical line at $\delta = 1$ in red in Figure 1. Second, from the expressions for the coefficients of the cubic, we can see that the sign of b (the coefficient on ν^2) changes sign at $\delta = 2$. Since roots of polynomials are impacted by signs of the coefficients, we make a further demarcation at $\delta = 2$. Thus, we have three bounded boxes to work with: *ABFE* (open on the right), *EFGH* (open on the left) and *GCDH* (closed on both sides).

On each of the three boxes, we can now use the condition in Proposition 1 to demarcate the *URR* (unique real root) area from the *NURR* (nonunique real root) area. This is depicted in Figure 1 by the green curves, which gives

the combination of δ and R_{max} at which $27q^2 + 4p^3 = 0$. The area above the green curve is the URR area and the area on and below the curve is the $NURR$ area.

To compute the bounding set for the ‘true’ treatment effect on any of the bounded boxes, we will create a vector, B_U , of size N^2 . The elements of this vector will be the roots of the cubic (15) computed at each of the N^2 points of an equally-spaced $N \times N$ grid that covers the URR area of the box. The vector B_U contains the set of possible values of treatment bias, ν . We can now use these values of ν to create compute $\beta^* = \tilde{\beta} - \nu$, the bias-adjusted treatment effect. Since $\beta^* \xrightarrow{p} \beta$ for *each* such ν , in large samples, each of these β^* are consistent estimates of the ‘true’ treatment effect. We then define a 95% bounding set for β as the interval defined by the 2.5-th and 97.5-th percentile of the empirical distribution of β^* . This interval would contain the ‘true’ treatment effect with 95% probability. To test whether the ‘true’ treatment effect is different from zero, we could see if this bounding set contains zero. We carry out the same procedure for each of the three boxes depicted in Figure 1.

3.3 Non-Unique Real Roots and Robustness

While the URR area allows us to generate credible bounding sets for the ‘true’ treatment effect, we also need to account for the area on the (δ, R_{max}) plane where the cubic (15) has multiple, i.e. three, real roots. To deal with this issue, let us identify the area of the (δ, R_{max}) plane where the cubic has multiple real roots.

Proposition 2. *If $27q^2 + 4p^3 \leq 0$, then the cubic equation in (15) will have three real roots.*

Proof. This is a well known result. Please see the appendix for some discussion. \square

The $NURR$ areas are depicted in each of the three boxes in Figure 1. Following the same steps as in the case of the URR area, we create a equally-spaced $N \times N$ grid that covers the $NURR$ area in each box. At each point of this grid, the cubic will have three real roots. Hence, we create three vectors, $B_{NU}^1, B_{NU}^2, B_{NU}^3$, each of size N^2 , and store the first, second and third real root in the vectors, respectively. We know from Proposition 2 in Oster (2019) that one of these vectors is the ‘true’ bias. To choose among the three vectors, I will pick the root whose empirical distribution is ‘closest’

to empirical distribution of the vector B_U , i.e. ‘close’ in terms of location.⁴ A simple way to choose the ‘closest’ distribution to that of B_U is to choose among B_{NU}^1, B_{NU}^2 and B_{NU}^3 the one whose median (or any other measure of central tendency) is nearest to the median (or any other measure of central tendency) of the empirical distribution of B_U . We proceed to compute the empirical distribution function of $\beta^* = \tilde{\beta} - \nu$ using the chosen root and compare that with the empirical distribution of $\beta^* = \tilde{\beta} - \nu$ computed with the vector B_U . If the two distributions are similar, then that provides robustness to our results. If they are dissimilar, then the method proposed here might be uninformative.

4 Keep Away from Equal Selection

The method proposed in the previous section to generate bounding sets for the ‘true’ treatment effect differs significantly from the method proposed by [Oster \(2019\)](#). In Oster’s proposal, a key role is played by the assumption of equal selection on observables and unobservables, i.e. $\delta = 1$. In this section, I show why this assumption leads to serious problems.

4.1 What is the Solution with Equal Solution?

What will be the solution for bias under equal selection? If we impose the restriction that $\delta = 1$ on the coefficients of the cubic in [\(15\)](#) we get,

$$\begin{aligned} a &= 0 \\ b &= -\tau_X \left(\overset{\circ}{\beta} - \tilde{\beta} \right) \sigma_X^2 \\ c &= \left(R_{max} - \tilde{R} \right) \sigma_Y^2 \left(\sigma_X^2 - \tau_X \right) - \left(\tilde{R} - \overset{\circ}{R} \right) \sigma_Y^2 \tau_X - \sigma_X^2 \tau_X \left(\overset{\circ}{\beta} - \tilde{\beta} \right)^2 \\ d &= \left(R_{max} - \tilde{R} \right) \sigma_Y^2 \left(\overset{\circ}{\beta} - \tilde{\beta} \right) \sigma_X^2, \end{aligned}$$

which converts the cubic in [\(15\)](#) to a quadratic equation in ν ,

$$b_1 \nu^2 + c_1 \nu + d_1 = 0, \tag{22}$$

⁴We are not concerned with the shape of the distributions.

where the coefficients of this quadratic are given by,

$$b_1 = -\tau_X \left(\overset{\circ}{\beta} - \tilde{\beta} \right) \sigma_X^2 \quad (23)$$

$$c_1 = \left(R_{max} - \tilde{R} \right) \sigma_Y^2 \left(\sigma_X^2 - \tau_X \right) - \left(\tilde{R} - \overset{\circ}{R} \right) \sigma_Y^2 \tau_X - \sigma_X^2 \tau_X \left(\overset{\circ}{\beta} - \tilde{\beta} \right)^2 \quad (24)$$

$$d_1 = \left(R_{max} - \tilde{R} \right) \sigma_Y^2 \left(\overset{\circ}{\beta} - \tilde{\beta} \right) \sigma_X^2. \quad (25)$$

The solutions of the quadratic in (22) are given by

$$\nu = \frac{-c_1 \pm \sqrt{c_1^2 - 4d_1b_1}}{2b_1},$$

which are noted in Corollary 1 in Oster (2019, pp. 193). Our first result is that the solution of the quadratic equation in (22) is always real.

Proposition 3. *The quadratic equation in (22) either has a unique real root or two distinct real roots. It does not have any complex roots.*

Proof. The proof follows by noting that the discriminant of this quadratic equation is non-negative, i.e. $c_1^2 - 4d_1b_1 \geq 0$, because $c_1^2 \geq 0$, and

$$\begin{aligned} -4d_1b_1 &= -4 \left\{ \left(R_{max} - \tilde{R} \right) \sigma_Y^2 \left(\overset{\circ}{\beta} - \tilde{\beta} \right) \sigma_X^2 \right\} \left\{ -\tau_X \left(\overset{\circ}{\beta} - \tilde{\beta} \right) \sigma_X^2 \right\} \\ &= 4 \left(R_{max} - \tilde{R} \right) \sigma_X^4 \sigma_Y^2 \tau_X \left(\overset{\circ}{\beta} - \tilde{\beta} \right)^2 \\ &\geq 0 \end{aligned}$$

where the last inequality follows because $R_{max} \geq \tilde{R}$. □

The implication of this result is that, in general, there will be two real roots of the quadratic equation in (22). Hence, in general, there will be two values of the bias in the treatment effect, and hence two values of the bias-adjusted treatment effect, when there is equal selection on observables and unobservables. Without further assumptions, it is not possible to arrive at a unique solution for the bias or the bias-adjusted treatment effect. This immediately leads to the following question: what conditions are necessary to give us an unique solution for the quadratic in (22)? The unique root will arise if and only if the discriminant of the quadratic equation is identically equal to zero. I now show that the discriminant can be zero only if we impose additional assumptions. These assumptions are difficult to justify on either theoretical or empirical grounds.

4.2 Condition for Unique Solution

For the quadratic equation in (22) to have a unique real solution, the discriminant must be zero, i.e.,

$$\left\{ \left(R_{max} - \tilde{R} \right) \sigma_Y^2 \left(\sigma_X^2 - \tau_X \right) - \left(\tilde{R} - \mathring{R} \right) \sigma_Y^2 \tau_X - \sigma_X^2 \tau_X \left(\mathring{\beta} - \tilde{\beta} \right)^2 \right\}^2 + 4 \left(R_{max} - \tilde{R} \right) \sigma_Y^4 \tau_X \left(\mathring{\beta} - \tilde{\beta} \right)^2 = 0. \quad (26)$$

Defining $Z = R_{max} - \tilde{R}$, we can write the above condition as a quadratic equation in Z ,

$$A^2 Z^2 + (2AB + 4C) Z + B^2 = 0, \quad (27)$$

where

$$A = \sigma_Y^2 \left(\sigma_X^2 - \tau_X \right) > 0 \quad (28)$$

$$B = - \left[\left(\tilde{R} - \mathring{R} \right) \sigma_X^2 \tau_X + \sigma_X^2 \tau_X \left(\mathring{\beta} - \tilde{\beta} \right)^2 \right] < 0 \quad (29)$$

$$C = \sigma_Y^2 \left(\mathring{\beta} - \tilde{\beta} \right) \sigma_X^2. \quad (30)$$

The two roots of (27) are given by

$$Z_1, Z_2 = \frac{-(2AB + 4C) \pm \sqrt{(2AB + 4C)^2 - 4A^2 B^2}}{2A^2}. \quad (31)$$

Note that the discriminant of the quadratic equation in (27) reduces to $16C^2 + 16ABC$. Since $B < 0$, it is possible, though not necessary, for the discriminant, $16C^2 + 16ABC$, to be negative.⁵ Hence, there are two cases to consider.

Case 1. If the discriminant is negative, then both the roots of (27), Z_1, Z_2 , are complex numbers. In this case, the uniqueness analysis falls through. This is because it is meaningless to entertain the possibility that $Z = R_{max} - \tilde{R}$ is a complex number. What does this mean? Since \tilde{R} is a known real number, this implies that there is no real value of R_{max} that would make the discriminant of the quadratic equation in (22) to be zero. Hence, in this case, there does not exist a unique magnitude of the bias in the treatment effect, ν , and hence, it is not possible to find a unique bias-adjusted treatment effect, β^* .

⁵Since $\tilde{R} \geq \mathring{R}$, the term in the square bracket in the definition of B is positive. Hence, $B < 0$.

Case 2. If the discriminant is nonnegative, then both the roots of (27), Z_1, Z_2 , are real. Thus, there exists real values of R_{max} which would give a unique value of the bias, and hence, the bias-adjusted treatment effect. But not all possible values of R_{max} are permissible. We know that R_{max} is never smaller than \mathring{R} . Hence, we need necessary conditions to ensure that the solutions of (27) are nonnegative. This is given in

Proposition 4. *If $\mathring{\beta} < \tilde{\beta}$, then then both roots of (27) are real. One of these roots will be nonnegative only if*

$$\begin{aligned} \sigma_Y^2 (\sigma_X^2 - \tau_X) \left[\left(\tilde{R} - \mathring{R} \right) \sigma_X^2 \tau_X + \sigma_X^2 \tau_X \left(\mathring{\beta} - \tilde{\beta} \right)^2 \right] \\ + 2\sigma_Y^2 \left(\mathring{\beta} - \tilde{\beta} \right) \sigma_X^2 \leq 0. \end{aligned} \quad (32)$$

Proof. To see the first part, note that if $\mathring{\beta} < \tilde{\beta}$, then $C < 0$. Hence $C^2 + ABC \geq 0$. Hence, the discriminant of (27) is nonnegative. To see the second part, note that since the denominator of the expression for the roots in (31) is always positive, the sign of the roots are the same as the sign of the numerator. If $2AB + 4C > 0$, then the numerator is negative because the expression within the square root in (31) is nonnegative and less than $2AB + 4C$. Hence, we have the following: $2AB + 4C > 0 \implies Z_1 < 0$ and $Z_2 < 0$. The contrapositive of this statement gives us: $Z_1 \geq 0$ or $Z_2 \geq 0 \implies 2AB + 4C \leq 0$. Hence, $2AB + 4C \leq 0$ is the necessary condition for at least one root being nonnegative. Plugging the expression for A, B and C , this becomes

$$\sigma_Y^2 (\sigma_X^2 - \tau_X) \left[\left(\tilde{R} - \mathring{R} \right) \sigma_X^2 \tau_X + \sigma_X^2 \tau_X \left(\mathring{\beta} - \tilde{\beta} \right)^2 \right] + 2\sigma_Y^2 \left(\mathring{\beta} - \tilde{\beta} \right) \sigma_X^2 \leq 0,$$

which is the expression in (32). \square

The above analysis has important implications. If $\delta = 1$, i.e. there is equal selection on observables and unobservables, and the condition in (32) is satisfied, then either $R_{max} - \tilde{R}$ will be given by the positive root of (27) when one of the two roots is negative, or $R_{max} - \tilde{R}$ will attain two positive values given by both the roots of (27) when both roots are positive. In either case, once we choose to impose the restriction that $\delta = 1$, then there will either be a unique value of R_{max} or two possible values of R_{max} that are permissible. The choice of $\delta = 1$ implies these *specific* values of R_{max} . Researchers are no longer at liberty to choose any other value of R_{max} . There is an additional angle to consider with regard to the analysis of bias

under the assumption of equal selection and this is highlighted in the next result.

Corollary 1. *If the estimate of the treatment effect declines with the addition of controls, i.e. if $\hat{\beta} > \tilde{\beta}$, then a meaningful bias-adjusted treatment effect cannot be computed under the assumption of equal selection.*

Proof. Note, first, that if $\hat{\beta} > \tilde{\beta}$, it is no longer guaranteed that (27) will have real roots. Consider, further, the necessary condition in (32) and note that if $\hat{\beta} > \tilde{\beta}$, then the condition cannot be satisfied. This is because the second term

$$2\sigma_Y^2 (\hat{\beta} - \tilde{\beta}) \sigma_X^2$$

is positive. Since $\sigma_X^2 > \tau_X$ and $\tilde{R} \geq \hat{R}$, the first term

$$\sigma_Y^2 (\sigma_X^2 - \tau_X) \left[(\tilde{R} - \hat{R}) \sigma_X^2 \tau_X + \sigma_X^2 \tau_X (\hat{\beta} - \tilde{\beta})^2 \right]$$

is always positive. Hence the expression on the left hand side of the condition in (32) is positive. An application of proposition 4 then shows that the quadratic equation in (27) cannot have meaningful roots. This, in turn, implies that the discriminant of the quadratic equation in (22) cannot be zero. This implies that the quadratic equation in (22) cannot have a unique root. Hence, a meaningful bias-adjusted treatment effect cannot be computed. \square

The implication of this corollary is that in cases where the treatment effect falls with the addition of controls, i.e. $\hat{\beta} > \tilde{\beta}$, the uniqueness of the solution is impossible. Thus, if for some research it is found that the treatment effect decreases with the addition of controls, then we can be sure that for this particular research a unique bias-adjusted treatment effect cannot be computed under the assumption of equal selection.

4.3 Equal Selection as Lower Bound

The second method proposed by Oster (2019) to deal with the problem of nonuniqueness is to choose a value of R_{max} and then use it to compute the value of δ that would make $\beta = 0$ ('true' treatment effect is zero). In Proposition 3, Oster (2019) demonstrates that such a δ can be uniquely computed. While the problem of nonuniqueness is thus solved, it creates its own problem of interpretation. What should researchers do with the value of δ thus computed? Following Altonji et al. (2005), the recommendation made by Oster (2019) is to use $\delta = 1$ as a lower bound. If the computed value

of δ is higher than unity, researchers are recommended to conclude that the problem of omitted variable bias is not serious; if the computed value of δ is lower than unity, then researchers need to be worried about the problem of omitted variable bias.⁶ How do we make sense of this recommendation?

Let us return to the expression for δ given in (11): $\delta = \delta_1/\alpha_1$. Thus, δ is the ratio of the coefficients in two linear projections: the numerator is the coefficient on the index of unobservables in a linear projection of the treatment variable on the index of unobservables; the denominator is the coefficient on the index of observables in a linear projection of the treatment variable on the index of observables. Altonji et al. (2005) and Oster (2019) have interpreted δ as the ratio of explanatory power of the unobservables to the corresponding power of the observables in explaining the variation in the treatment variable. Using this interpretation, they have argued that $\delta = 1$ is a sensible lower bound.

When a researcher finds that the value of $\delta = \delta^* > 1$, then that means that such a value of δ^* would be needed to wipe out the reported nonzero treatment effect. That is, if unobservables were δ^* times more important in explaining the variation in the treatment variable than the observables then the treatment effect, which was zero in reality, would be mistakenly reported as the nonzero $\tilde{\beta}$. In testing the importance of Catholic school attendance on student outcomes, when Altonji et al. (2005) find that the computed value of δ is significantly higher than unity, they draw the conclusion that the problem of omitted variable bias is not serious.

We find that selection on unobservables would need to be 3.55 times stronger than selection on observables in the case of high school graduation, which seems highly unlikely. It would have to be 1.43 times stronger to explain the entire college effect, which is also unlikely (Altonji et al., 2005, pp. 155).

Oster (2019) has echoed this idea several times in her paper. But is this justified? Can δ^* be interpreted in this fashion to draw conclusions about omitted variable bias? When δ^* is equal to 3.55, for instance, it does not tell us anything about the bias at other value of δ . The implicit understanding seems to be that if δ^* is larger than unity, then it must mean that an actual value of $\delta < 1$ would not nullify the reported estimate of the treatment effect. For instance, in the example quoted above from Altonji et al. (2005), if it is unlikely that selection on unobservables is “3.55 times stronger than

⁶This recommendation has been followed by researchers, e.g. see Jaschke and Keita (2021, Table 3).

selection on observables”, then the question to ask is this: what would the the magnitude of bias (of the treatment effect) be if the degree of selection on unobservables was lower? For instance, can we be sure that a value of $\delta < 1$ would not wipe out the reported treatment effect? Neither [Altonji et al. \(2005\)](#) nor [Oster \(2019\)](#) provide any answers to this question. In fact, they do not even pose this question.

In a similar manner, if a researcher finds that the computed value of δ^* is less than unity, she might then conclude that the problem of omitted variable bias is serious in her study. This is also a premature conclusion because we do not know anything about bias for other values of δ . Hence, using the computed value of δ^* to draw conclusions about the problem of omitted variable bias is an incomplete exercise, and as I show below, can lead to misleading results. A better strategy is to use a range of values of δ and draw conclusions only on the basis of roots computed over a *range* of values. I now turn to illustrating my method with actual data, and in doing so, I will point out how Oster’s methodology often give misleading results.

5 Two Applications

In this section, I report results of applying my method on two data sets.

5.1 Returns to Education

I used a simulated data set created by [Oster \(2019, section 4.1\)](#) from the NLSY-79 cohort to study the effect of years of education on log wages. The data set is constructed by regressing log wage on education, experience, sex and a set of eight demographic and family background variables: region of residence, race, marital status, mother’s education, father’s education, mother’s occupation, father’s occupation and number of siblings. Given these regressors, the fitted value from this regression gives the ‘true’ wage associated with different levels of education - conditional on demography and family background. If we were to regress the fitted values on these same set of regressors, we would get an R-squared of 1. In realistic research, regressions typically do not give perfect fit. Hence, an orthogonal error term is added to the fitted value from the above regression to make the true R-squared from the regression of 0.45 ([Oster, 2019, pp. 197](#)). The fitted value will serve as the outcome variable - let us call this the ‘constructed log wage’ - in two experiments I conduct to quantify omitted variable bias.

In the first experiment (Experiment 1), I regress constructed log wage on education, experience, sex and six of the demographic and family back-

ground variables. I drop mother’s education and father’s education from the set of regressors and know that, by construction, this will produce omitted variable bias in the estimated returns to education. In the second experiment (Experiment 2), I regress constructed log wage on education, experience, sex and mother’s education and father’s education. I drop the other six demographic and family background variables: region of residence, race, marital status, mother’s occupation, father’s occupation and number of siblings. In this case too, I know that, by construction, the estimated returns to education will have omitted variable bias. I will compare the results in both cases to the ‘true’ returns to education that I estimate by the regression of the constructed log wage on education, experience, sex and all the eight demographic and family background variables. The ‘true’ returns to education is estimated as 0.089 (with a standard error of 0.003). The R-squared in this regression is 0.373.

[Figure 2 about here]

5.1.1 Experiment 1

To begin the investigation of Experiment 1, I first estimate the treatment effect by regressing the constructed log wage on education, experience, sex and six of the demographic and family background variables (excluding mother’s education and father’s education). I find the estimate of $\tilde{\beta}$ to be 0.093 (standard error = 0.000, R-squared = 0.355). Thus, leaving out mother’s and father’s education produces an overestimate of the returns to education, as would be expected if education is positively correlated with mother’s and father’s education.

I report results of the bias analysis about Experiment 1 in Figure 2 and Table 1. In the top panel of Figure 2, I identify the unique real root (URR) and the nonunique real root (NURR) areas associated with the cubic equation (15) for the following three bounded boxes: the left figure uses what I will refer to as region 1, $0.50 \leq \delta < 0.99$ and $\tilde{R} \leq R_{max} \leq 1$; the middle figure uses region 2, $1.01 \leq \delta < 2.0$ and $\tilde{R} \leq R_{max} \leq 1$; the right figure uses region 3, $2.0 \leq \delta < 5.0$ and $\tilde{R} \leq R_{max} \leq 1$. In each of these plots, the colored area depicts the NURR area and the white area denotes the URR area. In the bottom panel, I plot the corresponding empirical distribution of the bias-adjusted treatment effect, $\beta^* = \tilde{\beta} - \nu$, where ν is the unique real root computed in the corresponding URR area.

In left hand side of panel A of Table 1, I report bounding sets for the ‘true’ treatment effect as the interval formed by the 2.5th and 97.5th per-

centile of the empirical distribution of the bias-adjusted treatment effect computed over the URR area (and plotted in the bottom panel of Figure 2). If we used region 1 for our computation, the bounding set would be $[0.009, 0.093]$; if we used region 2, the bounding set would be $[0.346, 0.441]$; if we used region 3, the bounding set would be $[0.268, 0.375]$. In this particular case, we know that the true treatment effect is 0.089 because we are using simulated data. Hence, we can see that using region 1 gives us the correct bounding set.

In research with observational data, the choice of the correct region would not be so straightforward. If researchers use the method presented in this paper, they have two options. First, they might be able to use their knowledge about the omitted variable and how it is likely to be related to the treatment variable to choose the correct region to use for computing bias. Since δ is the ratio of the impact of an index of omitted variables to the impact of an index of the included variables on the treatment variable, the question will boil down to the sign and magnitude of these effects. If institutional or historical knowledge of the issue under investigation allows a researcher to choose the relevant region, that would be the best option. If that is not possible, then the next best alternative is to look at the bounding sets for all meaningful regions and then draw conclusions. For instance, in panel A, we see that the bounding set for region 1, 2 and 3 all exclude zero. Hence, even if we did not have information to choose the correct region, we would still be able to say with lot of certainty that the true treatment effect is different from zero. In many realistic cases, this is all we might be interested in knowing, i.e. whether the true treatment effect is different from zero.

For robustness, we can turn to the right hand side of panel A of Table 1. Here, I report bounding sets using the *NURR* area. When we used region 1 of the *NURR* area for computing roots of the cubic, the bounding set for β contains zero. For the other two cases, the bounding sets do not contain zero. A researcher might then be justified in concluding that the true treatment effect is different from zero with lot of certainty (because 5 of the 6 bounding sets do not include zero).

[Table 1 about here]

Let us now compare the above methodology to the one proposed in Oster (2019), starting with the computation of the bounding set for the ‘true’ treatment effect. In the left panel of Table 2, I report the magnitude of bias when $\delta = 1$ (equal selection on observables and unobservables). Recall that

the magnitude of bias, under the assumption of equal selection, are the roots of the quadratic in (22). To solve for the roots of the quadratic, we need to choose values for δ and R_{max} . With equal selection, we have $\delta = 1$ and for R_{max} , I choose 20 equally-spaced values between $\tilde{R} + 0.01$ and 1. This is justified because we know that $1 \geq R_{max} > \tilde{R}$.

The results in the left panel of Table 2 show a striking result: the value of the discriminant of the quadratic is always positive. This means that the quadratic has two real roots for these values of R_{max} . These two roots are reported in columns 3 and 4, respectively. At the bottom of Table 2, I report the values of R_{max} which would ensure a unique real root. These values of R_{max} are -0.048 and -24.498, both meaningless quantities (because R-squared is always a fraction lying between 0 and 1). Thus, if a researcher were to use the assumption of equal selection, she would *not* be able to estimate a unique magnitude of bias or a unique bias-adjusted treatment effect - other than if she were to use meaningless values of R_{max} . This raises serious questions about the bias-adjusted treatment effect whose distribution is displayed in Figure 3(a) in Oster (2019). I will return to this question about uniqueness once again when I discuss bias-adjusted treatment effects in the investigation of maternal behaviour and child outcomes in the next sub-section.

[Table 2 about here]

Let us now turn to the second method proposed in Oster (2019): for a chosen value of R_{max} , computing the value of δ that is consistent with a ‘true’ treatment effect of zero. Once we impose the condition that the ‘true’ treatment effect is zero, i.e. $\beta = 0$, then Proposition 3 in Oster (2019) shows that there is a negatively-sloped functional relationship between R_{max} and δ . If the researcher were to choose any value of R_{max} , she would get the associated value of δ . For instance, if she chose $R_{max} = 0.45$, which we know to be the true R-squared, she would get $\delta^* = 3.374$. She would then conclude, rightly in this case, that the true treatment effect is different from zero.

There are two difficulties of this method. First, in research with observation data, there is no way to choose a sensible value of R_{max} . Instead of relying on a particular magnitude of R_{max} it seems better to work with a range, which is what my method does. Second, it is not clear how to interpret a given value of the computed δ^* so far as it can give us any information about the severity of the omitted variable bias problem. What does the value of $\delta = 3.374$ tell us about the problem of omitted variable

bias? It tells us that if $R_{max} = 0.45$ and $\delta = 3.374$, then $\beta = 0$. But it does not tell us anything about β for other values of δ and R_{max} , not even in the neighbourhood of this chosen (δ, R_{max}) point. Can we say anything about β if $R_{max} = 0.45$ and $\delta = 2$, or $R_{max} = 0.45$ and $\delta = 0.75$? We cannot. Hence, it seems that a more robust strategy should use information on a region rather than a single point, which is what my method does.

5.1.2 Experiment 2

For Experiment 2, I estimate the treatment effect by regressing the constructed log wage on education, experience, sex, mother’s education and father’s education (excluding the other six of the demographic and family background variables). I find the estimate of $\hat{\beta}$ to be 0.098 (standard error = 0.000, R-squared = 0.177). Thus, leaving out the other six demographic and family background variables, even when we include mother’s and father’s education, produces an overestimate of the returns to education. The magnitude of the overestimate is larger than in Experiment 1, where we had left out mother’s and father’s education (but had included these demographic variables).

[Figure 3 about here]

Results of the bias analysis for Experiment 2 are reported in Figure 3 and Table 1. Figure 3 presents plots of the URR and NURR areas and the empirical distribution of the bias-adjusted treatment effect, β^* . For all the three regions, the distribution of β^* is overwhelmingly to the right of zero. Turning to the bounding sets in panel B of Table 1, we see that other than in the case of region 1, the bounding sets do not contain zero. Taking the evidence in bottom panel of Figure 3 and panel B of Table 1, we might be justified in concluding the the ‘true’ treatment effect is different from zero.

To compare the above methodology with the one proposed in Oster (2019), we again start with the computation of the bounding set for the ‘true’ treatment effect. In the right panel of Table 2, I report the magnitude of bias that I get when I impose the restriction that $\delta = 1$ (equal selection on observables and unobservables). Just like in the case of Experiment 1, if a researcher were to use the assumption of equal selection, she would not be able to compute any meaningful bias-adjusted estimate of the ‘true’ returns to education - because there is no unique root. Turning to the second method proposed in Oster (2019), I have reported the value of δ^* at the bottom of panel B in Table 1: if we choose $R_{max} = 0.45$ (which we know

is the true R-squared), she would get $\delta^* = 0.736$. If the researcher followed Oster’s methodology, she would incorrectly conclude that the ‘true’ treatment effect is zero. Here we have a case where the method of computing δ^* gives misleading results.

5.2 Impact of Maternal Behaviour on Child Outcomes

Now I turn to an illustration of my method using observational data. The substantive issue under investigation in this example is the impact of maternal behaviour on child outcomes. In particular two child outcomes are studied: a child’s standardized IQ score and a child’s birth weight. In the study of child IQ, three treatment variables are used in turn: months of breastfeeding, any drinking of alcohol in pregnancy, and an indicator for being low birthweight and preterm. In studying child birthweight, two treatment variables are used, one by one: maternal smoking during pregnancy, and maternal drinking during pregnancy. The following control variables are used for both studies: child race, maternal age, maternal education, maternal income, maternal marital status. The question of interest is whether the treatment variables, each on their own, have any causal impact on the outcome variables.

5.2.1 Analysis of Bias and Bounding Sets

In Table 3, I present the estimates of the treatment effect from the short and intermediate regressions. These results replicate the corresponding results in Table 3 in Oster (2019). For instance, if we look at the first row of Table 3, we see that the effect of (months of) breastfeeding on child IQ is 0.044 (column 1) in the short regression and 0.017 (column 3) in the intermediate regression. Moving from the short to the intermediate regression, the R-squared increases from 0.008 (column 2) to 0.249 (column 4). We can read all the other numbers in columns 1 through 4 in a similar manner. Since these models are likely to have omitted variables, we would like to quantify the effect of the omitted variable bias.

[Figure 3 about here]

I begin the analysis of bias by constructing the relevant cubic equation and identifying the *URR* and *NURR* areas and plotting the empirical distribution of the bias-adjusted treatment effect, β^* , that arises from solving the cubic equation over the *URR* area. Figure 4 presents the results about the regions and the empirical distribution plots of β^* corresponding the first

three rows in Table 3, i.e. in cases where the outcome variable is a child’s standardized IQ score (IQ). For the figure in the top panel, the treatment is months of breastfeeding; in the middle panel, the treatment is drinking during pregnancy; and in the bottom panel, the treatment is low birth weight and premature birth.

The figures can be interpreted exactly as we did in Figure 2 and 3: it depicts the *URR*/*NURR* areas, and below the colored area plots, we have a plot of the empirical distribution of the bias-adjusted treatment effect, β^* . One thing that requires comment are the spots that are missing figures. If we look at the plot for region 2 ($0.01 < \delta \leq 0.99$) in the middle panel, we see that there is no *URR* area (the whole area is violet). This means that a unique real root of the cubic equation cannot be computed in this case in the bounded box of region 2. Hence, $\beta^* = \tilde{\beta} - \nu$ cannot be computed, because there is no ν to use. That is why there is no empirical distribution plot of β^* . The same comment applies to region 2 ($1.01 \leq \delta \leq 2.00$) and region 3 ($2.00 \leq \delta \leq 5.00$) in the bottom panel of Figure 4.

Figure 5 presents results about the regions and the empirical distribution of β^* corresponding to row 4 and 5 in Table 3, i.e. where the outcome variable is birth weight of the child (BW). The top panel has smoking during pregnancy as the treatment variable; in the bottom panel, drinking during pregnancy is the treatment variable. For each analysis, we have nonzero *URR* area. Hence, we do not have any missing plots in Figure 5.

We can turn to Table 4 to find estimates of bounding sets corresponding to the relevant empirical distributions of β^* . The top panel of Table 4 presents bounding sets for the ‘true’ effect of breastfeeding on a child’s standardized IQ score. If we consider the *URR* area, region 2 and region 3 give bounding sets that do not include zero, the former being $[0.301, 0.318]$ and the latter being $[0.275, 0.304]$, but region 1’s bounding set, $[-0.038, 0.017]$, does include zero.

In observational studies like these, we do not have the luxury of knowing which is the correct region to use. What we can say, therefore, on the basis of this analysis, is only a conditional statement: if relative selection on unobservables, δ , lies between 0 and 1, then the ‘true’ treatment effect might be zero; if, on the other hand, the relative selection on unobservables, δ , is larger than unity, then the ‘true’ treatment effect is certainly different from zero. In the latter case, the treatment effect estimated from the intermediate regression, 0.017 (row 1, column 3, Table 3), is unlikely to be wiped out even after we take account of possible omitted variable bias. Our conclusion would not change if we also used the *NURR* area (right side of top panel in Table 4).

The panels in Table 4 present results corresponding to those same rows in Table 3. Thus the second panel in Table 4 corresponds to the second row in Table 3, and so on. For rows where bounding sets are missing, it is the case that the area of the URR region is zero. Hence, no unique real roots exist. Thus, the bias-adjusted treatment effect could not be computed. Hence, there is no bounding set. This means that for these cases, the method proposed in this paper does not provide any guidance about the magnitude of bias.

[Figure 4 about here]

5.2.2 Oster’s Identified Set is Misleading

The conclusions one would draw from the use of my method will be very different from those that would emerge from using Oster’s method. One of the key differences relate to the bounding set (what Oster calls the ‘identified set’). Oster has reported bounding sets for the ‘true’ treatment effect in column 5 in Table 3 in Oster (2019). To compute the bounding set according to Oster’s methodology, one would need to find the roots of the quadratic in (22) for a chosen value of δ and R_{max} . Following Oster (2019), I choose $\delta = 1$ and $R_{max} = 0.61$ – because these are the parameter values used in panel A in Table 3 in Oster (2019). For these parameter values, I get positive values of the discriminant. Hence, one cannot get a unique value of the root for these parameter values and so, one cannot compute a unique magnitude of $\beta^*(\delta = 1, R_{max} = 0.61)$. This means that a unique identified set cannot be computed. To highlight this difficulty, I report the values of R_{max} that would be necessary to make the discriminant 0 - the condition that is stated in (26). As can be seen from columns 5 and 6 in Table 3, the values of R_{max} are largely meaningless. Only in row 4, do the values of R_{max} have meaningful magnitudes. But, then, there is no reason to choose these particular values of R_{max} for computing the bias. Hence, even these meaningful magnitudes in row 4 are difficult to justify. The conclusion must be that the ‘identified set’ reported in column 5 in Table 3 in Oster (2019) are not meaningful.

If the quadratic in (22) does not have a unique root for $\delta = 1$ and $R_{max} = 0.61$, then how can Oster (2019, column 5, Table 3) compute an identified set? She does not explain this in the paper and so I can only venture a guess. It seems that she has taken recourse to Assumption 3 to generate a unique root. Assumption 3, in Oster (2019), states that the sign of the covariance between the treatment variable and the actual index of

observables is the same as the sign of the covariance between the treatment variable and the predicted index of observables. The meaning and import of this assumption is explained thus.

Effectively, this assumes that the bias from the unobservables is not so large that it biases the *direction* of the covariance between the observable index and the treatment. Under Assumption 3, if $\delta = 1$, there is a unique solution (Oster, 2019, pp. 194, emphasis in original).

It is not clear how the sign restriction on the covariance between the treatment variable and the index of observables can generate a unique root of the quadratic equation in (22). Oster (2019) does not provide a proof of this important claim. If we turn to the condition for a unique root given in (26), we see that all variables in that equation, other than R_{max} , are given. This implies that the necessary and sufficient condition for a unique root is a specific magnitude of R_{max} - the magnitude that solves (26). A sign restriction on the covariance between the treatment and the index of observables cannot guarantee its existence.⁷

5.3 Does δ^* Provide Useful Information?

Let us consider the second strategy proposed in Oster (2019), i.e. computing and using δ^* (the relative selection on unobservables that is consistent with a zero treatment effect). In column 7 in Table 4, I have reported the values of δ^* . The proposal in Oster (2019) is to use a value of $\delta^* < 1$ as evidence that omitted variable bias is serious, and a value of $\delta^* > 1$ as evidence that the problem of omitted variable bias can be ignored. The bounding sets I have computed, and presented in Table 4, suggest that such conclusions might be misleading.

In row 1, Table 3, the value of δ^* is 0.366. If we followed Oster’s methodology, we would conclude that the reported estimate of the treatment effect (0.017) is not reliable. Once we take account of omitted variable bias, the true treatment effect is likely to be zero - this is what Oster’s method would suggest. If we turn to the bounding sets reported in the top panel of Table 4, we see that this conclusion is not wholly warranted. We see that

⁷In the STATA code that implements her method, she uses the condition that sign of $\hat{\beta} - \beta$ must be the same as the sign of $\beta - \beta$ to choose one of the roots. Quite apart from the fact that this is an *ad hoc* assumption, even this will not ensure that the condition in (26) holds. Hence, even this *ad hoc* assumption, whose relationship to Assumption 3 is itself not clear, will not guarantee a unique solution. Thus, the STATA code is incorrectly picking up one of the solutions and artificially solving the problem of nonuniqueness.

if $\delta > 1$, the bounding set will not include zero. For instance, if one (or several) of the omitted variables is such that it is strongly correlated with the treatment variable (breastfeeding), then the relative degree of selection on unobservables might very well be larger than unity. In that case, the conclusion drawn on the basis of $\delta^* = 0.366$ would be incorrect.⁸

Let us contrast the above analysis with the analysis of row 3 in Table 3. In this case, the treatment variable is LBW+preterm (low birth weight and preterm birth). From column 7, we see that $\delta^* = 1.36$. If we followed Oster’s methodology, we would conclude that the problem of omitted variable bias is not serious. Turning to the third panel in Table 4, we see that such a conclusion might not be warranted. To see why, let us think about how we interpret the finding that $\delta^* = 1.36$. We interpret this to mean only a high degree of selection on unobservables, i.e. $\delta = 1.36$, could completely nullify the reported treatment effect (-172.51) and make it zero. Underlying such an argument is the implicit understanding that a lower degree of selection on unobservables would not cause any problem. The results in the third panel in Table 4 shows that if $0.01 \leq \delta \leq 0.99$, the bounding set for the ‘true’ treatment effect contains zero. Thus, even with a lower degree of selection on unobservables, it is possible that the true treatment effect vanish. Hence, if we use the computed value of $\delta^* = 1.36$ to conclude that the problem of omitted variable bias is not serious, we would be drawing an incorrect conclusion.

6 Concluding Comments

Omitted variable bias is an ubiquitous problem in applied econometric work. Quantifying the magnitude of bias and computing bias-adjusted treatment effects is an important area of research. Building on earlier work by [Altonji et al. \(2005\)](#), in a recent contribution, [Oster \(2019\)](#) has proposed a novel methodology to compute bias-adjusted treatment effect when there is proportional selection on observables and unobservables. In this paper, I have argued that while [Oster \(2019\)](#) posed the problem correctly, her proposed solutions are problematic. I have instead proposed an alternative methodology to compute bounding sets for the true treatment effect. To conclude the discussion, let me give a quick summary of the proposed methodology for the benefit of applied researchers.

- Estimate the short regression by regressing the outcome variable on

⁸Recall that in Experiment 2, we have already seen a case where a $\delta^* < 1$ led to an incorrect conclusion.

the treatment variable (and exclude all controls). Store the coefficient as $\hat{\beta}$ and the R-squared as \hat{R} .

- Estimate the intermediate regression by regressing the outcome variable on the treatment variable and all the controls. Store the coefficient as $\tilde{\beta}$ and the R-squared as \tilde{R} .
- Estimate an auxiliary regression by regressing the treatment variable on all the controls. Store the variance of the residual as τ_X .
- Store the variance of the outcome variable as σ_y^2 and the variance of the control variable as σ_X^2 .
- Form the cubic equation in (15).
- Choose δ_{min} and δ_{max} and construct the three bounded boxes depicted in Figure 1. Identify the *URR* (unique real root) area and the *NURR* areas in each box using the condition in Proposition 1.
- For each box, choose a $N \times N$ grid to cover the *URR* area and solve the cubic at each point on the grid. Collect the $N^2 \times 1$ vector of real roots, ν , of the cubic equation. This gives the empirical distribution of the treatment bias.
- Define $\beta^* = \tilde{\beta} - \nu$ and find the empirical distribution of the bias-adjusted treatment effect. Use the empirical distribution of β^* to define bounding sets for the ‘true’ treatment effect.

The method outlined above will give bounding sets for three different ranges of the magnitude of selection, $0 < \delta < 1$, $1 < \delta < 2$, and $2 < \delta < 5$. If zero is not contained in any of these three bounding sets, then the researcher can conclude with lot of confidence that the true treatment effect is different from zero. If, on the other hand, zero is contained in each of the three bounding sets, then the researcher would be sure that the problem of omitted variable is serious enough to completely nullify whatever results she has found.

In many cases, this happy outcome will not occur and researchers will face ambiguity. This will arise when some of the bounding sets will and some will not include zero. In such cases, a researcher will have to draw on knowledge of the institutional details of the substantive issue under investigation in identifying a correct range of δ . For instance, if a particular research question has an omitted variable that is understood to be very important,

then it might be justified to use $\delta > 1$. If, on the other hand, the researcher is sure that all important variables have been included in the model, and hence, that the omitted variable is relatively less important, then a range of $\delta < 1$ might be justified.

I would like to end by pointing out a weakness of the method proposed in this paper. The main drawback is that the method is conservative. The bounds that are produced by this method are not sharp, they are expansive. Of course this cannot be avoided unless we have more information about the true value of δ . If we knew the true value of δ we would have been able to produce a very sharp bounding set for the true β . Since we lack knowledge about the true value of δ , we have to experiment with meaningful ranges of values. No wonder, the bounding sets produced by this method are very wide and conservative. This suggests that one fruitful direction for future research to extend this methodology is to think of ways of forming relatively precise estimates of δ (the degree of selection on unobservables). The more precise values of δ we generate, the tighter bounds we can generate for the true treatment effect, β .

References

- Altonji, J. G., Elder, T. E., and Taber, C. R. (2000). Selection on Observed and Unobserved Variables: Assessing the Effectiveness of Catholic Schools. Working Paper 7831, National Bureau of Economic Research, Cambridge, MA.
- Altonji, J. G., Elder, T. E., and Taber, C. R. (2005). An evaluation of instrumental variable strategies for estimating the effects of catholic schooling. *The Journal of Human Resources*, 40(4):791–821.
- Basu, D. (2020). Bias of OLS Estimators due to Exclusion of Relevant Variables and Inclusion of Irrelevant Variables. *Oxford Bulletin of Economics and Statistics*, 82(1):209–234.
- Galor, O. and Ozak, O. (2016). The Agricultural Origins of Time Preference. *American Economic Review*, 106(10):2586–2624.
- Greene, W. H. (2012). *Econometric Analysis*. Prentice Hall.
- Hellesland, J., Challamel, N., Casandjian, C., and Lanos, C. (2013). *Reinforced Concrete Beams, Columns and Frames: Section and Slender Member Analysis: Mechanics and Design*. ISTE Ltd. and John Wiley & Sons, Inc, London, UK.

- Jaschke, P. and Keita, S. (2021). Say it like Goethe: Language learning facilities abroad and the self-selection of immigrants. *Journal of Development Economics*, 149:102597.
- Michalopoulos, S. and Papaioannou, E. (2016). The Long-Run Effects of the Scramble for Africa. *American Economic Review*, 106(7):1802–1848.
- Oster, E. (2019). Unobservable Selection and Coefficient Stability. *Journal of Business and Economic Statistics*, 37(2):187–204.
- Wooldridge, J. M. (2002). *Econometric Analysis of Cross Section and Panel Data*. MIT Press.
- Goldsmith-Pinkham, P., Sorkin, I., and Swift, H. (2019). Bartik Instruments: What, When, Why, and How. *American Economic Review*, 110(8):2586–2624.

Table 1: Estimates of Bounding Set of the ‘True’ Treatment Effect in Experiment 1 and Experiment 2

	URR Region		NURR Region	
	Lower Bound	Upper Bound	Lower Bound	Upper Bound
	(1)	(2)	(3)	(4)
<u>Panel A: Experiment 1</u>				
Region 1	0.009	0.093	-0.035	0.089
Region 2	0.346	0.441	0.419	0.812
Region 3	0.268	0.375	0.326	0.432
Memo: $\delta^*=3.374$				
<u>Panel B: Experiment 2</u>				
Region 1	-0.011	0.097	-0.305	0.080
Region 2	0.252	0.279	0.262	0.318
Region 3	0.246	0.267	0.264	0.281
Memo: $\delta^*=0.736$				

Notes: This table reports the bounding set for the true treatment effect in Experiment 1 and Experiment 2, both of which estimate the returns to education by regressing log wage on education and a set of controls (for details see section 5.1). In Experiment 1, I include 6 demographic variables but drop mother’s education and father’s education from the baseline wage regression; in Experiment 2, I include mother’s education and father’s education but drop the other 6 demographic variables from the baseline wage regression. The value of δ^* is computed according to Proposition 3 in Oster (2019) using $R_{max} = 0.45$. URR region = unique real root region; NURR region = nonunique real root region. See figure 1 for details.

Table 2: Estimates of Bias with Equal Selection

Experiment 1				Experiment 2			
(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
R_{max}	$Disc$	$Root_1$	$Root_2$	R_{max}	$Disc$	$Root_1$	$Root_2$
0.365	0.019	-0.961	0.001	0.187	0.001	-0.236	0.005
0.399	0.018	-0.936	0.005	0.230	0.001	-0.224	0.026
0.432	0.017	-0.910	0.010	0.273	0.001	-0.214	0.049
0.465	0.016	-0.886	0.015	0.315	0.001	-0.206	0.074
0.499	0.016	-0.861	0.020	0.358	0.001	-0.199	0.101
0.532	0.015	-0.837	0.025	0.401	0.001	-0.194	0.128
0.566	0.014	-0.813	0.030	0.444	0.002	-0.189	0.156
0.599	0.014	-0.789	0.036	0.487	0.002	-0.185	0.185
0.633	0.013	-0.766	0.042	0.529	0.002	-0.181	0.215
0.666	0.013	-0.743	0.049	0.572	0.002	-0.178	0.245
0.699	0.012	-0.720	0.056	0.615	0.003	-0.176	0.276
0.733	0.012	-0.698	0.063	0.658	0.003	-0.174	0.306
0.766	0.011	-0.676	0.071	0.701	0.003	-0.172	0.337
0.800	0.011	-0.655	0.079	0.743	0.004	-0.170	0.369
0.833	0.011	-0.634	0.088	0.786	0.004	-0.168	0.400
0.866	0.010	-0.614	0.097	0.829	0.005	-0.167	0.432
0.900	0.010	-0.594	0.107	0.872	0.005	-0.166	0.464
0.933	0.010	-0.575	0.118	0.914	0.006	-0.164	0.496
0.967	0.009	-0.556	0.129	0.957	0.006	-0.163	0.528
1	0.009	-0.538	0.140	1	0.007	-0.163	0.560

Memo: R_{max} necessary for $Disc=0$

Experiment 1: -0.048 -24.498

Experiment 2: -0.002 -42.152

Notes: This table reports the magnitude of the discriminant ($Disc$) and roots of the quadratic equation ($Root_1, Root_2$) that defines treatment bias under equal selection (equation 26). The roots are computed for 20 values of R_{max} in the interval $\hat{R} + 0.01 \leq R_{max} \leq 1$. For details of Experiment 1 and Experiment 2, see section 5.1 and the footnote to Table 1.

Table 3: Estimates of Bias Using Oster's Method

	(1)	(2)	(3)	(4)	(5)	(6)	(7)
	$\hat{\beta}$	\hat{R}	$\tilde{\beta}$	\tilde{R}	R_{max}	R_{max}	δ^*
<u>Dep Var: IQ</u>							
Breastfeed	0.044	0.045	0.017	0.256	-0.02	-159.127	0.366
Drink in Preg	0.176	0.008	0.05	0.249	-0.589	-96.432	0.26
LBW + Preterm	-0.188	0.004	-0.125	0.251	1132.733	4.346	1.36
<u>Dep Var: Birthweight</u>							
Smoking in Preg	-183.115	0.319	-172.51	0.352	0.378	0.319	0.821
Drink in Preg	-16.668	0.301	-14.149	0.338	1.997	1.341	0.667

Notes: This table replicates some of the results reported in Table 3 in Oster (2019). Other than columns (5) and (6), all results are reported in Table 3 in Oster (2019). The reported value of δ^* in row 4 is different from the corresponding magnitude in Table 3 in Oster (2019). The values of R_{max} reported in columns (5) and (6) correspond to values that would be needed to ensure a unique real root for the quadratic in (22).

Table 4: Bounding Set for True Treatment Effect in the Study of the Impact of Maternal Behaviour on Child Outcomes

	URR Region		NURR Region	
	Lower Bound	Upper Bound	Lower Bound	Upper Bound
	(1)	(2)	(3)	(4)
IQ~Breastfeed:R1	-0.038	0.017	-0.08	0.014
IQ~Breastfeed:R2	0.301	0.318	0.315	0.379
IQ~Breastfeed:R3	0.275	0.304	0.283	0.316
IQ~Drink in Preg:R1	-0.218	0.048	-0.332	0.038
IQ~Drink in Preg:R2	NA	NA	NA	NA
IQ~Drink in Preg:R3	4.788	5.316	4.87	5.811
IQ~LBW+Preterm:R1	-0.124	0.017	-0.119	0.056
IQ~LBW+Preterm:R2	NA	NA	NA	NA
IQ~LBW+Preterm:R3	NA	NA	NA	NA
BW~Smoking in Preg:R1	-171.086	525.835	-1873.559	-442.629
BW~Smoking in Preg:R2	-884.723	-344.41	-166.342	519.094
BW~Smoking in Preg:R3	-634.045	-310.65	-162.711	13.65
BW~Drink in Preg:R1	-13.842	37.712	51.143	59.743
BW~Drink in Preg:R2	-582.457	-477.48	-2350.762	-580.96
BW~Drink in Preg:R3	-481.99	-224.473	-671.087	-385.222

Notes: This table reports the bounds for the ‘true’ treatment effect of maternal behaviour on child outcomes. Each panel in this table corresponds to the same row in Table 3. The estimated treatment effects are reported in Table 3 and in Oster (2019, Table 3). IQ=standardized IQ score; BW=birthweight. For further details see section 5.2. To estimate bias, the roots of the cubic in (15) are computed over the following regions: R1= $\{0.01 \leq \delta \leq 0.99; \tilde{R} \leq R_{max} \leq 1\}$; R2= $\{1.01 \leq \delta \leq 2.00; \tilde{R} \leq R_{max} \leq 1\}$; R3= $\{2.01 \leq \delta \leq 5.00; \tilde{R} \leq R_{max} \leq 1\}$.




Figure 1: *The construction of a bounded box defined by $\delta_{\min} \leq \delta \leq \delta_{\max}$ and $\tilde{R} \leq R_{\max} \leq 1$, and its demarcation into the URR (unique real root) and NURR (nonunique real root) area by the equation denoting the discriminant of a cubic equation being zero (depicted by the green curves). The bounded box is trifurcated by $\delta_{\min} \leq \delta < 1$, $1 < \delta \leq 2$ and $2 \leq \delta \leq \delta_{\max}$.*




Figure 2: The figure identifies the region of unique real root (top panel) and plots the histogram and empirical density function (red solid line) of β^* (bottom panel) for Experiment 1 (for details see section 5.1). In the top panel, the white region identifies all the possible combinations of (δ, R_{max}) such that the cubic equation defining the treatment effect has a unique real root; the violet region identifies the area where a unique real root does not exist. In Experiment 1, I include 6 demographic and family background variables but drop mother's education and father's education from the baseline wage regression.




Figure 3: The figure identifies the region of unique real root (top panel) and plots the histogram and empirical density function (red solid line) of β^* (bottom panel) for Experiment 2 (for details see section 5.1). In the top panel, the white region identifies all the possible combinations of (δ, R_{\max}) such that the cubic equation defining the treatment effect has a unique real root; the violet region identifies the area where a unique real root does not exist. In Experiment 2, I include mother's education and father's education but drop the other 6 demographic and family background variables from the baseline wage regression.




Figure 4: The figure identifies the region of unique real root (top panel) and plots the histogram and empirical density function (red solid line) of β^* (bottom panel) for impact of maternal behaviour on IQ score of the child (for details, see section 5.2). In the top panel, the treatment is months of breast-feeding, in the middle panel, the treatment is drinking during pregnancy, and in the bottom panel, the treatment is low birth weight and premature. The color coding is the same as in Figure 2.




Figure 5: *The figure identifies the region of unique real root (top panel) and plots the histogram and empirical density function (red solid line) of β^* (bottom panel) for impact of maternal behaviour on birthweight of the child (for details, see section 5.2). In the top panel, the treatment is smoking during pregnancy, and in the bottom panel, the treatment is drinking during pregnancy. The color coding is the same as in Figure 2.*

A Unique Real Root of a Cubic Equation

Solving cubic equations is common in the engineering literature and for this presentation I draw partly on [Hellesland et al. \(2013, Appendix 1\)](#). Consider the cubic equation in t ,

$$at^3 + bt^2 + ct + d = 0, \quad (33)$$

where $a \neq 0$. Divide through by a to get

$$t^3 + \frac{b}{a}t^2 + \frac{c}{a}t + \frac{d}{a} = 0. \quad (34)$$

A change of variable,

$$x = t + \frac{b}{3a},$$

can convert this into a ‘depressed’ cubic,

$$x^3 + px + q = 0, \quad (35)$$

where,

$$p = \frac{3ac - b^2}{3a^2}$$
$$q = \frac{27a^2d + 2b^3 - 9abc}{27a^3}.$$

To solve (35), we will express x as the difference of two numbers, i.e. $x = a - b$. Since,

$$(a - b)^3 + 3ab(a - b) - (a^3 - b^3) = 0, \quad (36)$$

we will get back (35) from (36), where $x = a - b$, if the following two conditions are satisfied:

$$ab = \frac{p}{3} \quad (37)$$

and

$$a^3 - b^3 = -q. \quad (38)$$

Thus, if we are able to solve for a and b in terms of p and q , we will be able to get $x = a - b$, and from that we will be able to finally get the value of $t = x - (b/3a)$.

Note that the above two conditions, (37) and (38), show that the sum and product of a^3 and $(-b)^3$ are $-q$ and $-p^3/27$, respectively. But this

means that a^3 and $(-b)^3$ are the roots of the following quadratic equation in y ,

$$y^2 + qy - \frac{p^3}{27} = 0. \quad (39)$$

Denoting one of the roots of the quadratic as a^3 , we have,

$$a^3 = -\frac{q}{2} + \sqrt{\frac{q^2}{4} + \frac{p^3}{27}} = U_1$$

and denoting the other root as $-b^3$, we get

$$-b^3 = -\frac{q}{2} - \sqrt{\frac{q^2}{4} + \frac{p^3}{27}},$$

so that

$$b^3 = \frac{q}{2} + \sqrt{\frac{q^2}{4} + \frac{p^3}{27}} = U_2.$$

The solutions of the cubic equation (33) will depend on the sign of the discriminant

$$D_3 = \frac{q^2}{4} + \frac{p^3}{27} = \frac{27q^2 + 4p^3}{108}. \quad (40)$$

Proposition 5. *If $D_3 > 0$, then the cubic equation has one real root and two complex roots. The unique real root is given by*

$$t_1 = \sqrt[3]{-\frac{q}{2} + \sqrt{\frac{q^2}{4} + \frac{p^3}{27}}} - \sqrt[3]{\frac{q}{2} + \sqrt{\frac{q^2}{4} + \frac{p^3}{27}}} - \frac{b}{3a},$$

and the complex roots are given by

$$t_2 = \omega \sqrt[3]{-\frac{q}{2} + \sqrt{\frac{q^2}{4} + \frac{p^3}{27}}} - \omega^2 \sqrt[3]{\frac{q}{2} + \sqrt{\frac{q^2}{4} + \frac{p^3}{27}}} - \frac{b}{3a},$$

and

$$t_3 = \omega^2 \sqrt[3]{-\frac{q}{2} + \sqrt{\frac{q^2}{4} + \frac{p^3}{27}}} - \omega \sqrt[3]{\frac{q}{2} + \sqrt{\frac{q^2}{4} + \frac{p^3}{27}}} - \frac{b}{3a},$$

where ω is the cube root of unity given by

$$\omega = e^{i\frac{2\pi}{3}} = -\frac{1}{2} + \frac{\sqrt{2}}{3}i$$

and $i = \sqrt{-1}$.

Proof. To see this, note that the possible pairs of (a, b) that will satisfy (37) and (38) are

$$\left(\sqrt[3]{U_1}, \sqrt[3]{U_2}\right), \left(\omega \sqrt[3]{U_1}, \omega^2 \sqrt[3]{U_2}\right), \left(\omega^2 \sqrt[3]{U_1}, \omega \sqrt[3]{U_2}\right).$$

Since $D_3 > 0$, U_1 and U_2 are real numbers. Hence, the unique real value of x is given by

$$x = a - b = \sqrt[3]{U_1} - \sqrt[3]{U_2} = \sqrt[3]{-\frac{q}{2} + \sqrt{\frac{q^2}{4} + \frac{p^3}{27}}} - \sqrt[3]{\frac{q}{2} + \sqrt{\frac{q^2}{4} + \frac{p^3}{27}}}$$

and the corresponding unique real root of the original cubic equation (33) is given by

$$t = \sqrt[3]{-\frac{q}{2} + \sqrt{\frac{q^2}{4} + \frac{p^3}{27}}} - \sqrt[3]{\frac{q}{2} + \sqrt{\frac{q^2}{4} + \frac{p^3}{27}}} - \frac{b}{3a}.$$

The other two roots will be complex conjugate numbers because they involve ω . \square

An immediate corollary follows. Cubic equations with real coefficients can have either one or three real roots. Complex roots occur in conjugate pairs. Thus, when the cubic has only real roots, i.e. three real roots, it will be the case that $D_3 \leq 0$.