Unobservable Selection and Coefficient Stability: Theory and Evidence*

Emily Oster

Brown University and NBER

January 26, 2015

Abstract

A common heuristic for evaluating robustness of results to omitted variable bias is to observe coefficient movements after inclusion of controls. This heuristic is informative only if selection on observables is informative about selection on unobservables. I formalize this link through a proportional selection relationship, following Altonji, Elder and Taber (2005). I show that it is necessary to take into account coefficient movements *and* movements in R-squared values in identifying omitted variable bias. I further demonstrate that in the empirically common case with multiple observed controls it is also necessary to account for the share of the variation in treatment accounted for by control variables. I describe a formal bounding argument for omitted variable bias under the proportional selection relationship. I show two validation exercises suggesting that this performs well empirically. I discuss application of this procedure to a large set of publications in economics, and use evidence from randomized studies to draw guidelines for bounding values.

*Ling Zhong, Unika Shrestha, Damian Kozbur, Guillaume Pouliot, David Birke and Angela Li provided excellent research assistance. I thank David Cesarini, Raj Chetty, Todd Elder, Amy Finkelstein, Guido Imbens, Larry Katz, Matt Gentzkow, Matt Notowidigdo, Chad Syverson, Manisha Shah, Azeem Shaikh, Jesse Shapiro, Bryce Steinberg, Matt Taddy, Heidi Williams and participants in seminars at Brown University, University of Chicago Booth School of Business, Wharton and Yale for helpful comments. I am grateful to a number of authors for providing replication files or re-running analysis by request. I gratefully acknowledge financial support from the Neubauer Family. *Stata* code to perform the calculations described in this paper is available from the authors website or through ssc under the name **psacalc.**

1 Introduction

Concerns about omitted variable bias are common to most or all non-experimental work in economics.[1] The most straightforward approach to such concerns is to include controls that can be observed. Angrist and Pischke (2010) argue that among the major advances in empirical economics in the past two decades is greater effort to identify the most important threats to validity, and to address them with appropriate selection of controls. Even with careful selection of controls, however, the possibility of bias from unobserved controls remains.

A common heuristic for evaluating the robustness of a result to omitted variable bias concerns is to look at the sensitivity of the treatment effect to inclusion of observed controls. In three top general interest economics journals in 2012, 75% of non-experimental empirical papers included such sensitivity analysis.[2] The intuitive appeal of this approach lies in the idea that the bias arising from the observed controls is informative about the bias that arises from the unobserved ones. This is not, however, implied by the baseline assumptions underlying the linear model.

Formally, using the observables to identify the bias from the unobservables requires making further assumptions about the covariance properties of the two sets. On one extreme, if the unobservables are completely unlike the observables, nothing about the remaining bias is learned from the inclusion of the observables. The case in which the bias is fully identified is the other extreme: where one assumes *all* of the unobservables share the same covariance properties as the observables (Murphy and Topel, 1990; Altonji, Elder and Taber, 2005; Altonji et al, 2011).

Even under this most optimistic assumption, however, coefficient movements alone are not a sufficient statistic to calculate bias. To illustrate why, consider the case of a researcher estimating wage returns to education with individual ability as the only confound, and where there are two orthogonal components of ability, one of which has a higher variance than the other. Assume wages would be fully explained if both ability components were observed but, in practice, the researcher sees only one of the two. The coefficient will appear much more stable if the observed ability control is the lower variance one, but this is not because the bias is smaller but simply because less of the wage outcome is explained.

This example is described in more detail in Section 2. The key observation is that the quality of the control will be diagnosed by the movement in R-squared when the control is included. This observation generalizes to *all* cases in which the observed controls share covariance properties with the unobserved

[1]Despite recent trends, this still makes up the vast majority of results within economics: in 2012 the combination of the *American Economic Review,* the *Quarterly Journal of Economics* and the *Journal of Political Economy* published 69 empirical, non-structural papers, only 11 of which were randomized.

[2]The sample includes non-structural papers in the *American Economic Review, Journal of Political Economy* and *Quarterly Journal of Economics.* The link between coefficient stability and omitted variable bias is often quite direct. For example, Chiappori et al (2012) state: "It is reassuring that the estimates are very similar in the standard and the augmented specifications, indicating that our results are unlikely to be driven by omitted variables bias." Similarly, Lacetera et al (2012) state: "These controls do not change the coefficient estimates meaningfully, and the stability of the estimates from columns 4 through 7 suggests that controlling for the model and age of the car accounts for most of the relevant selection."

controls. Omitted variable bias is proportional to coefficient movements, but only if such movements are scaled by movements in R-squared. This point is closely related to the partial R-squared logic in Imbens (2003).

The importance of R-squared movements is rarely acknowledged in discussions of coefficient stability. In Section 2 I demonstrate this using a sample of 57 top journal papers in economics that present coefficient stability evidence. Only 5 mention anything about R-squared movements and 30% do not even report R-squared values. Moreover, there is little relationship between the change in coefficient and the change in R-squared, suggesting that coefficient movements alone do not diagnose bias. As a further point, the basic underlying assumption of some proportional selection relationship, which is required for coefficient movements to have any relation with bias, is typically neither acknowledged nor tested.

This paper discusses how coefficient stability formally relates to robustness to omitted variable bias. I make explicit the link between coefficient movements, R-squared movements and omitted variable bias through the assumption of related covariance. I suggest a structured notion of robustness. I perform two validation examples – one based on constructed data and one linking possibly biased observational relationships to external causal estimates – which suggest this procedure performs well. Finally, I apply this adjustment to a set of papers in economics and use insights from randomized data to suggest standards for robustness.

I begin in Section 3 with the theory. The theory draws heavily on the setup in Altonji, Elder and Taber (2005) and their related papers; I focus on the linear model case. To facilitate intuition, I first develop the estimator under the assumption that there is a single observable control and equal selection on observed and unobserved variables. Formally, I consider the following model: $Y = \beta X + \gamma_1 w_1^o + W_2 + \epsilon$, where W_2 is unobserved and therefore cannot be included in the estimation.[3] I impose the equal selection assumption: $\frac{Cov(\gamma_1 w_1^o, X)}{Var(\gamma_1 w_1^o)} = \frac{Cov(W_2, X)}{Var(W_2)}$. Under this assumption, the omitted variable bias on the treatment effect estimated with controls can be expressed as an intuitive function of the coefficient and R-squared values. In particular, the residual omitted variable bias after inclusion of controls is proportional to the coefficient movements and the ratio of the movement in R-squared with inclusion of the observable control to the expected movement in R-squared with the inclusion of the unobservable controls.

I then outline the general estimator for the bias when there are multiple observed controls. In most empirical settings researchers include multiple observed controls in their regressions, so this generalization is a crucial one for empirical applications. I relax the assumption of equal selection and replace it with a *proportional selection relationship* in which the covariance relationship above is proportional, not necessarily equal, and relies on a coefficient of proportionality, δ. The estimator in this case can be expressed as a function of δ, the coefficient and R-squared movements, as well as information on the variance of the treatment and the share of that variance that is explained by the observed controls. In many cases, the simple estimator

[3] As I detail more later, the approach in this paper starts from the idea that there is a full model, defined as stated, and that β is an object of interest. There are a closely related set of issues about whether this model is the correct one; if it is not, then β will not be a "causal" effect (see, e.g., Leamer (1978)). Throughout I will refer to β as a "treatment effect" with the understanding that this is done in the context of this model.

developed first will provide a close approximation to the general estimator, suggesting intuition about the bias can be easily developed from coefficient and R-squared movements. However, this is not always true. Importantly, it is possible in the case with multiple controls that the coefficient will appear stable (in fact, may be completely unchanged) and there may still be large bias on the estimated treatment effect.

A key input into either estimator is the R-squared from a hypothetical regression of the outcome on treatment and both observed and unobserved controls; I denote this $R_{max.}$ If the outcome can be fully explained by the treatment and full controls set, then $R_{max} = 1$. In many empirical settings it seems likely (due, for example, to measurement error) that the outcome cannot be fully explained even if the full control set is included. Knowledge about measurement error or expected idiosyncratic variation in the outcome can be used to develop intuition about this value.

Following the theory I discuss implementation, with the goal of suggesting a formal expression of robustness that might replace heuristic statements about small coefficient movements. As in Altonji, Elder and Taber (2005) I suggest that equal selection (i.e. $\delta = 1$) may be an appropriate upper bound on δ. Essentially, this argues that the unobservables should not be more important than the observables in explaining the treatment. I then suggest that researchers adopt a bounding value for R_{max} – ideally a conservative upper bound, denoted $\overline{R_{max}}$ – and report either the value of δ for which the estimator would produce a treatment effect of zero or the value of β which is produced by $\delta = 1$ and $\overline{R_{max}}$. In the former case, a value of $\delta > 1$ would be seen as suggesting a robust result; in the latter, showing that the adjusted β leads to the same conclusion would be a natural standard. The latter essentially argues for the construction of an "identified set," akin to the logic in Tamer (2010) and Manski (2003), bounded on one side by the controlled treatment effect and on the other by the bias-adjusted effect with $\delta = 1$ and $R_{max} = \overline{R_{max}}$.

Following the theory, I turn to the primary contribution of the paper, which is to test the performance of this estimator in data and evaluate the use of coefficient stability in papers within economics.

Section 5 first uses NLSY data to construct a dataset relating education and wages; the data is constructed such that we know the true treatment effect. I evaluate the performance of this adjustment by excluding combinations of controls from the "observed" set. I estimate the value of δ which would be produced by each excluded set and calculate the bias-adjusted treatment effect. I show in 89% of cases using the bounding logic described in above would produce a set which includes the true effect; in only 62% of cases does the confidence interval of the naive controlled coefficient include this effect. This may actually undervalue this performance as the control set is selected here at random rather than based on using the most important controls first, as would be common in practice. I show in this case that approximating the general estimator with the simple estimator would make little difference to the conclusions.

In a second test I estimate several relationships between maternal behavior and child outcomes; socioeconomic status confounds are a major concern. I match possibly biased observational estimates with external evidence on causal effects from randomized data or comprehensive meta-analyses (this is close in spirit to LaLonde (1986)). I then ask whether the robustness tests described above would separate true from false associations. I find that the adjustment performs well: the approach identifies as robust only the two relationships for which external evidence confirms a link. I find that in this case the simple approximation would perform less well and substantially understates the bias in some relationships.

Both of the validation exercises suggest empirical support for this assumption.

In the final section of the paper I turn to the application of this procedure to the economics literature. I focus on two questions: (1) How do stability statements in published papers in economics hold up to a version of this adjustment?; (2) Is it possible to make a general statement about bounds on R_{max}?

I begin with a sample of papers in the *American Economic Review, Journal of Political Economy, Quarterly Journal of Economics* and *Econometrica,* published between 2008 and 2013 and satisfying a set of citation cutoffs. I extract all relationships for which a coefficient stability heuristic is reported (57 papers; 131 results). I limit to cases where it is possible to access replication files, which is necessary for calculating some of the inputs; this limits the sample to 89 results. I calculate the bias-adjusted treatment effect with $\delta = 1$ and varying bounding values for $R_{max.}$ My primary definition of robustness is whether this exercise rejects zero; I also explore an auxiliary definition related to coefficient size.

Only about 20% of results are robust to a value of $R_{max} = 1$. I show other bounds on R_{max} which are a function of the fully controlled R-squared. These capture the idea that there is variation in how predictable outcomes are, and this variation can be roughly inferred from how much is predicted by the observables. Denoting the fully controlled R-squared as $\tilde{R}$, I explore robustness to $R_{max} = \Pi\tilde{R}$, with varying values of Π. About 37% of results are robust to a value of $\Pi = 2$, and 66% to a value of $\Pi = 1.25$.

I compare the conclusions from the general estimator to the simple approximation. Eighty percent of the time the simple approximation would lead to the same conclusions, although the error is sometimes sizable. In 90% of cases the simple approximation understates the bias, on average by around 30%. This points to the importance of performing the full bias calculation.

There is considerable variation across papers in the robustness of these stability claims, but this does not suggest an appropriate general value for the bound on R_{max}. For that, I turn to randomized results. The claim that the coefficient is unchanged by inclusion of controls implicitly suggests that the treatment is assigned as if randomly. If that is the case, then the coefficient movement should be within the bounds we would see if treatment were randomized. It is common in randomized papers to show coefficients with and without controls, either as a balancing test or to increase precision.

I draw a sample of all randomized papers from the *American Economic Review, Journal of Political Economy, Quarterly Journal of Economics*, *Econometrica* and *American Economic Journal: Applied Economics* between 2008 and 2013 which report coefficients with and without controls and for which I can access replication files (65 results). I derive cutoffs based on values of Π which would allow 90% of randomized results to survive: this value is $\Pi = 1.3$.

This provides a full robustness reporting standard. I suggest that researchers either (1) report the value of δ for which $\beta = 0$ with $R_{max} = 1.3\tilde{R}$ and show it exceeds 1 or; (2) calculate the bias-adjusted β with $\delta = 1$ and $R_{max} = 1.3\tilde{R}$ and show it leads to the same conclusion.

In the full sample of non-randomized results considered, about 54% would survive this bounding robustness argument. I conclude this section by discussing some examples which illustrate the importance of both taking into account the movement in R-squared and consider the full estimator rather than just the sample approximation.

This paper adds to a large literature on causal inference in the face of unobserved confounds (Rosenbaum and Rubin, 1983). Imbens (2003) presents an analysis of sensitivity using a partial R-squared logic which is conceptually similar to the insights here. A number of methodological papers consider the approach of varying the covariate set as a sensitivity analysis (Heckman and Hotz, 1989; Dehejia and Wahba, 1999; Gelbach, *forthcoming*). I also relate to the idea of specification and control set search (Leamer, 1978; Pearl, 2000; Angrist and Pischke, 2010). It is worth noting that the approach here differs in some conceptual sense from the latter set of references in that I am concerned with estimating a bias-adjusted treatment effect under an assumption about the full model rather than in searching for the appropriate full model.

In the formal use of the proportional selection relationship I follow several recent papers (Murphy and Topel, 1990; Altonji, Elder and Taber, 2005; Altonji et al, 2011). The primary contribution to this literature is in connecting more explicitly with the intuitive methodologies used by many empirical researchers in evaluating bias, and by connecting the theory directly to empirical work.

The rest of the paper is organized as follows. Section 2 provides an illustrative example of the issues raised above and describes the use of the coefficient stability heuristic in economies. Section 3 describes the theory and Section 4 briefly discusses implementation. Section 5 performs the validation tests and Section 6 turns to applications within economics. Section 7 concludes.

2 Coefficient Stability Heuristic: Illustrative Example and Use in Economics

I motivate the analysis in the paper with a simple illustration of the issues here, and with some data on coefficient stability within economics.

Illustrative Example

A central point of this paper is to make clear that coefficient movements alone are not sufficient to discuss bias, even under the strong assumption of related observed and unobserved variables. As an illustration, consider the case of a researcher estimating wage returns to education with individual ability as the only confound.[4] Assume that ability has two orthogonal components, each of which contribute to determining wages. Further, assume that wages would be fully explained if an exact measure of ability could be observed but, in practice, the researcher sees only one of the two ability controls. Finally, assume that one of the components of ability has a significantly larger variance than the other so, as a result, accounts for a larger share of the variance in wages.

Consider this setup under the central assumption of the paper - namely, that education relates to the observed and unobserved components of ability equally, so a regression of education on either ability control yields the same coefficient. The key observation is that if the researcher observes the ability control with the lower variance the coefficient will appear stable when the control is included. This is not, however, because the bias is small, but simply because the control is less important in explaining wages.

To see this precisely, consider Panel A of Table 1. This panel uses constructed data in which the true treatment effect is zero and there are two orthogonal confounds, one of which has a larger variance than the other. The first row shows controlled and uncontrolled coefficients when the observed control is the one with the larger variance; the second shows the coefficients when the observed control is the one with the smaller variance. The coefficient in the second row appears much more stable, even though the true effect is zero in both.

The key difference in the two rows is the change in R-squared, which diagnoses the poor quality of the proxy in the second row compared to the first. The uninformative control leaves the coefficient largely unchanged but also adds little to the R-squared. This observation generalizes to *all* cases in which the observed controls share covariance properties with the unobserved controls. Omitted variable bias is proportional to coefficient movements, but only if such movements are scaled by movements in R-squared.

The converse of this point in made in Panel B. Here, I consider two constructed examples in which the coefficient movement is identical but the movements in R-squared vary widely. In the first row, the small coefficient movement is accompanied by a large move in R-squared; in the second row, the move in R-squared is very small. To the extent that we would like to draw conclusions about the true treatment effect from the controlled coefficient, our intuition suggests that we will come closer in the first case than the second. The fact that so much of the outcome is explained in the controlled regression suggests that there is simply very little variation left to bias the coefficient.

[4] This example is motivated by independent work by Pischke and Schwandt (2013), although their setting is focused on issues of measurement error in the ability measure.

Indeed, if we assume that in both cases the observed and unobserved variables would together explain all of the variation in Y and these two sets relate to the treatment X in the same way (this is the equal selection assumption discussed much more below) the final column of Panel B shows the implied true treatment effect. Clearly it is only in the first row that reasonable conclusions could be drawn based on the controlled coefficient.

Coefficient Stability in Economics

The discussion above makes clear the importance of incorporating movements in R-squared in coefficient stability discussions. In empirical work in economics, however, the importance of the R-squared movements is rarely acknowledged in these discussions.

To elaborate on this claim, I extract all papers in the *American Economic Review, Quarterly Journal of Economics*, *The Journal of Political Economy* and *Econometrica* from 2008-2010 with at least 20 citations in the ISI Web of Science, and those from 2011-2013 in the same journals with at least 10 citations. From these papers I extract all results where the researcher explores the sensitivity of the result to a control set. This sample (full citation list in Appendix D) includes 57 papers with 131 total results. Only 5 of these papers mention anything about R-squared movements, and 29% of them do not report the R-squared values in the paper.

In principle, if coefficients and R-squared values typically move together, it is possible this omission would not meaningfully affect conclusions. That is, if large coefficient movements are always accompanied by large R-squared movements, then the coefficient stability is effectively a sufficient statistic. Similarly, if the controlled R-squared values are always very large – say, always close to 1 – then the coefficient movements would be enough. In practice, neither of these is the case.

Figure 1 uses the results extracted from the 57 papers described above.[5] I limit the sample to results where the controlled effect is significant. The figures graph the relationship between the percent movement in effect size and the absolute movement in R-squared values. Figure 1a uses all results, and Figure 1b limits to cases where the inclusion of controls moves the coefficient toward zero.

It is not the case that the controlled regressions uniformly have a high R-squared. The range of values for the controlled R-squared is 0.0005 to 0.9894, with an average of 0.42. Moreover, there is at best a very weak relationship between coefficient movements and R-squared movements. If we limit to results where the percent change in coefficient values is between -12% and -8%, the range of changes in R-squared values is from 0.008 to 0.29.

To develop one concrete comparison, we can consider points (A) and (B) in Figure 1b. In the result in

[5]In the 29% of cases where R-squared was not reported I use replication files to estimate or request these from the researcher. In 2 cases it was not possible to obtain R-squared values.

point (A), the coefficient decreases from 0.49 to 0.43 with inclusion of controls, and the R-squared increases from 0.44 to 0.93. In the result in point (B), the coefficient change is in a similar range – from 0.21 to 0.19 – but the increase in R-squared is only 0.36 to 0.37. These sets of figures could have very different implications for the true treatment effect. If we assume that in both cases the observed and unobserved variables would together explain all of the variation in Y and these two sets relate to the treatment X in the same way, the true treatment effect for point (A) is 0.42, whereas the true effect for point (B) is -1.23.

The remainder of this paper develops an estimator for bias under proportional selection. This will provide a formal way to relate coefficient and R-squared movements to omitted variable bias. This formal development will also lay bare a second issue. When there are multiple controls included (as is common in most applications) coefficient stability may be misleading even in the presence of sizable R-squared movements. This provides a stronger argument for presenting formal results on bias adjustment rather than relying on heuristic statements.

In the final section of the paper I will return to the economics literature and revisit these data and results in light of the estimator developed. It is worth noting that the theory here draws on Altonji, Elder and Taber (2005) who present a similar methodology for evaluating robustness. They do not connect directly to the idea of coefficient movements and, perhaps as a result, the empirical literature has been slow to adopt this approach. The papers cited above all post-date their work but in only one paper is any formal bias calculation done and, as I detail more in Section 6, it is not done correctly.

3 Theory

I begin in this section by developing the simple case of a single observable variable and equal selection on observed and unobserved variables. The solution in this case is intuitive. The second subsection derives and discusses the general estimator, including a brief discussion of inference. The third subsection discusses the relation between the two estimators.

Before moving into the formal theory it is useful to briefly discuss the conceptual approach here. I focus on on an approach to estimating an unbiased treatment effect from a model in which there are some observed confounders and some unobserved confounders. I will refer to the "full model" as a model which controls for both the observed and unobserved variables. The assumptions on relative selection will relate the treatment to the observed and unobserved confounders in this model. The bias-adjusted treatment effect estimated will be a causal effect *if* the full model here is one from which one can derive a causal estimate. There are numerous other threats to causality, including poorly specified functional form and others.[6] It is possible there are a set of unobserved controls which *do not* share covariance properties with the observed controls. This approach

[6]The analysis here is appropriate only for a linear model. Altonji, Elder and Taber (2005) develop a similar estimator in the context of at least one type of non-linear model (a bivariate probit model).

will not remove the bias arising from those. Leamer (1978) and Angrist and Pischke (2010) (among others) provide more discussion on the related topic of model specification in search of a causal effect.

The value of the approach here is that it addresses at least one known threat to causality. It is one which, based on the evidence in Section 2, much empirical work appears to focus on.

3.1 Single Observable, Equal Selection

Consider the regression model

$$Y = \beta X + \gamma_1 w_1^o + W_2 + \epsilon \tag{1}$$

X represents the treatment and the coefficient of interest is β; w_1^o and W_2 represent confounders. Specifically, w_1^o is an observed control variable with true coefficient γ_1. W_2 is a vector which is a linear combination of *unobserved* control variables w_j^u, multiplied by their true coefficients: $W_2 = \sum_{j=1}^{J_u} w_j^u \gamma_j^u$. Define $W_1 = \gamma_1 w_1^o$. W_1 is therefore an index of the observed control multiplied by its true coefficient. I assume that ϵ is orthogonal to X, w_1^o and W_2.

This setup is drawn from Altonji, Elder and Taber (2005) who provide a different, also consistent, estimator in this case. In their case, they assume $\epsilon = 0$ and that W_2 contains some error unrelated to X. I argue this alternative formulation will be useful in moving to empirical work. However, the results here go through in a straightforward way if $\epsilon = 0$.

I assume that $Cov(W_1, W_2) = 0$ and that $Var(X) = \sigma_{XX}$. The assumption of orthogonality between W_1 and W_2 is discussed in more detail below. The covariance matrix associated with the vector $[X, W_1, W_2]'$ is positive definite. Note that without further assumptions on the relationship between X, w_1^o and W_2 there is no information provided about the bias associated with W_2 by seeing the bias from w_1^o.

Define the equal selection relationship as $\frac{\sigma_{1X}}{\sigma_{11}} = \frac{\sigma_{2X}}{\sigma_{22}}$, where $\sigma_{iX} = Cov(W_i, X)$, $\sigma_{ii} = Var(W_i)$. Note this relationship is defined on the index W_1, not directly on the variable w_1^o.

Define the coefficient resulting from the short regression of Y on X as $\mathring{\beta}$ and the R-squared from that regression as $\mathring{R}$. Define the coefficient from the intermediate regression of Y on X and w_1^o as $\tilde{\beta}$ and the R-squared as $\tilde{R}$. Finally, define R_{max} as the R-squared from a hypothetical regression of Y on X, w_1^o and W_2. Note these are in-sample values.

The omitted variable bias on $\mathring{\beta}$ and $\tilde{\beta}$ is controlled by the auxiliary regressions of (1) w_1^o on X; (2) W_2 on X; and (3) W_2 on X and w_1^o. Denote the in-sample coefficient on X from regressions of w_1^o and W_2 on X as $\hat{\lambda}_{w_1^o|X}$ and $\hat{\lambda}_{W_2|X}$, respectively and the coefficient on X from a regression of W_2 on X and w_1^o as $\hat{\lambda}_{W_2|X,w_1^o}$. Denote the population analogs of these values $\lambda_{w_1^o|X}$, $\lambda_{W_2|X}$ and $\lambda_{W_2|X,w_1^o}$.

All estimates are implicitly indexed by n. Probability limits are taken as n approaches infinity. All observations are independent and identically distributed according to model (1). By standard omitted variable bias formulas, I can express the probability limits of the short and intermediate regression coefficients in terms of these values:

$$\begin{aligned}\mathring{\beta} &\xrightarrow{p} \beta + \gamma_1 \lambda_{w_1^o|X} + \lambda_{W_2|X} \\ \tilde{\beta} &\xrightarrow{p} \beta + \lambda_{W_2|X,w_1^o}\end{aligned}$$

The asymptotic bias on $\tilde{\beta}$ (the coefficient on X with controls included) is $\lambda_{W_2|X,w_1^o}$ which, given the definitions above, is equal to $\frac{\sigma_{22}\sigma_{1X}}{\sigma_{11}(\sigma_{XX}-\frac{\sigma_{1X}^2}{\sigma_{11}})}$. Denote this bias as Π.

Define the following.

$$\beta^* = \tilde{\beta} - \left[\mathring{\beta} - \tilde{\beta}\right] \frac{R_{max} - \tilde{R}}{\tilde{R} - \mathring{R}}$$

Proposition 1 summarizes the result.

Proposition 1. $\beta^* \xrightarrow{p} \beta$.

Proof. I outline the proof here, with details in Appendix A.1. Using the definition of coefficient and R-squared values and recalling the bias is denoted Π we have the following relationships.

$$\begin{aligned}(\mathring{\beta} - \tilde{\beta}) &\xrightarrow{p} \left(\frac{\sigma_{1X}}{\sigma_{XX}}\right)\left(1 - \frac{\sigma_{1X}}{\sigma_{11}}\Pi\right) \\ \left(\tilde{R} - \mathring{R}\right)\hat{\sigma}_{yy} &\xrightarrow{p} \sigma_{11} + \Pi^2(\sigma_{XX} - \frac{\sigma_{1X}^2}{\sigma_{11}}) - \frac{1}{\sigma_{XX}}\left(\sigma_{1X} + \Pi(\sigma_{XX} - \frac{\sigma_{1X}^2}{\sigma_{11}})\right)^2 \\ (R_{max} - \tilde{R})\hat{\sigma}_{yy} &\xrightarrow{p} \Pi\left(\frac{\sigma_{11}\left(\sigma_{XX} - \frac{\sigma_{1X}^2}{\sigma_{11}}\right)}{\sigma_{1X}} - \Pi\left(\sigma_{XX} - \frac{\sigma_{1X}^2}{\sigma_{11}}\right)\right)\end{aligned}$$

These define a system of three equations in three unknowns (σ_{11},σ_{1X} and Π). The system is identified and the solution is $\Pi = \left[\mathring{\beta} - \tilde{\beta}\right] \frac{R_{max}-\tilde{R}}{\tilde{R}-\mathring{R}}$. □

Some intuition for this result may be developed by observing that $\Pi = \tilde{\beta} - \beta$ so this result implies that $\frac{\tilde{\beta}-\beta}{\mathring{\beta}-\tilde{\beta}} = \frac{R_{max}-\tilde{R}}{\tilde{R}-\mathring{R}}$. That is, under the equal selection assumption, the ratio of the movement in coefficients is equal to the ratio of the movement in R-squared. The objects W_1 and W_2 enter the equation for Y symmetrically in terms of coefficients, and equal selection implies they also are symmetric in their impact on X. The only way in which their impact may differ is if they have different variances. This possible difference will be captured in the differential contributions to R-squared. In the special case where the variances are equal, then $\frac{R_{max}-\tilde{R}}{\tilde{R}-\mathring{R}} = 1$ and the coefficient movement with inclusion of observed controls is equal to the expected coefficient movement with unobserved controls.[7]

It is important to note that the setup and result here is exactly what we would derive if we were able to

[7] This is the special case dealt with in Bellows and Miguel (2009).

observe W_1 directly, rather than w_1^o. This is because the coefficient and R-squared values from the short and intermediate regressions are the same in either case.

Proportional Selection

Define the proportional selection relationship as $\delta\frac{\sigma_{1X}}{\sigma_{11}} = \frac{\sigma_{2X}}{\sigma_{22}}$, where δ is the coefficient of proportionality. Equal selection corresponds to the case of $\delta = 1$. With $\delta \neq 1$ the estimator $\beta^* = \tilde{\beta} - \delta\left[\mathring{\beta} - \tilde{\beta}\right]\frac{R_{max}-\tilde{R}}{\tilde{R}-\mathring{R}}$ will be a close approximation for the bias as long as δ is relatively close to 1. This observation leads to the possibility of calculating an approximate value for $\hat{\delta}$ which is the degree of selection for which β is equal to some target value $\hat{\beta}$; when $\hat{\beta} = 0$ this tells us the degree of selection on unobservables relative to observables which would be sufficient to explain away the result. Specifically, $\hat{\delta} \approx \frac{(\tilde{\beta}-\hat{\beta})(\tilde{R}-\mathring{R})}{(\mathring{\beta}-\tilde{\beta})(R_{max}-\tilde{R})}$.

For the case where $\hat{\beta} = 0$ this is equivalent to the estimator developed in Altonji, Elder and Taber (2005) and is consistent under the null that $\beta = 0$.

3.2 General Estimator

I consider now the general case where selection is proportional and there are possibly multiple variables in the set of observable controls. I retain much of the notation from above and the proof method proceeds similarly.

Consider the regression model

$$Y = \beta X + \Psi\boldsymbol{\omega}^{\mathbf{o}} + W_2 + \epsilon \tag{2}$$

where $\boldsymbol{\omega}^{\mathbf{o}}$ is a vector of the observed controls, $\omega_1^o...\omega_j^o$. The index W_2 is not observed. Define $W_1 = \Psi\boldsymbol{\omega}^{\mathbf{o}}$ and assume that all elements of $\boldsymbol{\omega}^{\mathbf{o}}$ are orthogonal to W_2, so W_1 and W_2 are orthogonal. Without loss of generality, assume the elements of $\boldsymbol{\omega}^{\mathbf{o}}$are also orthogonal to each other.[8] Define the proportional selection relationship as $\delta\frac{\sigma_{1X}}{\sigma_{11}} = \frac{\sigma_{2X}}{\sigma_{22}}$, where $\sigma_{iX} = Cov(W_i, X)$, $\sigma_{ii} = Var(W_i)$ and δ is the coefficient of proportionality. Note that at this point we do not make any assumptions about δ so this relationship will always hold for some δ.

The orthogonality of W_1 and W_2 is central to deriving the results here and may be somewhat at odds with the intuition that the observables and the unobservables are "related." In practice, the weight of this assumption is in how we think about the proportionality condition. To see that, consider the case where the elements of $\boldsymbol{\omega}^{\mathbf{o}}$ are correlated with W_2. The coefficient of proportionality is some δ_1. Now define $\tilde{W_2}$ as the residual from a regression of W_2 on $\boldsymbol{\omega}^{\mathbf{o}}$. By definition $\tilde{W_2}$ is orthogonal to $\boldsymbol{\omega}^{\mathbf{o}}$ and all the coefficients and R-squared values defined below will be identical to the original setup. The coefficients on the elements of $\mathbf{w}^{\mathbf{o}}$ will differ, but these do not factor into the calculations below. The only difference will be the use of a new degree of proportionality, $\delta_2 \neq \delta_1$.

Denote the coefficient resulting from the short regression of Y on X as $\mathring{\beta}$ and the R-squared from that

[8] All results go through identically if these elements are correlated.

regression as $\mathring{R}$. Define the coefficient from the intermediate regression of Y on X and $\omega^{\mathbf{o}}$ as $\tilde{\beta}$ and the R-squared as $\tilde{R}$. Finally, define R_{max} as the R-squared from a hypothetical regression of Y on X, $\omega^{\mathbf{o}}$ and W_2. Note these are in-sample values.

It is worth noting that in this case even though there are multiple controls I consider the approach of adding them all at once rather than one at at time. Gelbach (*forthcoming*) develops some theory and intuition around the process of sequentially adding controls in a similar setup.

The omitted variable bias on $\mathring{\beta}$ and $\tilde{\beta}$ is controlled by the auxiliary regressions of (1) each value $\omega_1^o..\omega_j^o$ on X; (2) W_2 on X; and (3) W_2 on X and $\omega^{\mathbf{o}}$. Denote the in-sample coefficient on X from regressions of each ω_i^o on X as $\hat{\lambda}_{\omega_i^o|X}$ and the in-sample coefficient on X from a regression of W_2 on X as $\hat{\lambda}_{W_2|X}$. Finally, denote the coefficient on X from a regression of W_2 on X and $\omega^{\mathbf{o}}$ as $\hat{\lambda}_{W_2|X,\omega^{\mathbf{o}}}$. Denote the population analogs of these values $\lambda_{\omega_i^o|X}$, $\lambda_{W_2|X}$ and $\lambda_{W_2|X,\omega^{\mathbf{o}}}$.

Define $Var(X) = \sigma_{XX}$. Define $\tilde{X}$ as the residual from a regression of X on $\omega^{\mathbf{o}}$. Define the variance of this residual in sample as $\hat{\tau}_x$ and the population analog as τ_x. Denote the sample variance of Y as $\hat{\sigma}_{yy}$ and note that $\hat{\sigma}_{yy} \xrightarrow{p} \sigma_{yy}$.

All estimates are implicitly indexed by n. Probability limits are taken as n approaches infinity. All observations are independent and identically distributed. As above, I can express the probability limits of the short and intermediate regression coefficients in terms of these values:

$$\begin{aligned}
\mathring{\beta} &\xrightarrow{p} \beta + \sum_{i=1}^{j} \psi_i^o \lambda_{\omega_i^o|X} + \lambda_{W_2|X} \\
\tilde{\beta} &\xrightarrow{p} \beta + \lambda_{W_2|X,\omega^{\mathbf{o}}}.
\end{aligned}$$

Under the proportional selection relationship, the asymptotic bias on $\tilde{\beta}$ is $\frac{\delta \sigma_{1X} \sigma_{22}}{\sigma_{11} \tau_x}$. Denote this bias Π.

Define the cubic function $f(\nu)$ as:

$$\begin{aligned}
f(\nu) &= \delta\left((R_{max} - \tilde{R})\sigma_{yy}\right)\left(\mathring{\beta} - \tilde{\beta}\right)\sigma_{XX} \\
&+ \nu\left(\delta\left((R_{max} - \tilde{R})\sigma_{yy}\right)(\sigma_{XX} - \tau_x) - \left(\left(\tilde{R} - \mathring{R}\right)\sigma_{yy}\right)\tau_x - \sigma_{XX}\tau_x\left(\mathring{\beta} - \tilde{\beta}\right)^2\right) \\
&+ \nu^2\left(\tau_x A\left(\mathring{\beta} - \tilde{\beta}\right)\sigma_{XX}(\delta - 2)\right) \\
&+ \nu^3(\delta - 1)(\tau_x \sigma_{XX} - \tau_x^2)
\end{aligned}$$

Proposition 2. *The proposition has two cases depending on the roots of $f(\nu)$.*

Case 1: *$f(\nu)$ has a single real root, define this root as ν_1. Define $\beta^* = \tilde{\beta} - \nu_1$. $\beta^* \xrightarrow{p} \beta$.*

Case 2: *$f(\nu)$ has three real roots, define them as ν_1, ν_2 and ν_3. Define a set $\beta^* = \{\tilde{\beta} - \nu_1, \tilde{\beta} - \nu_2, \tilde{\beta} - \nu_3\}$. One element of the set β^* converges in probability to β.*

Proof. I outline the proof here, with details in Appendix A.2. Using the definition of coefficient and R-squared values and recalling the bias is denoted Π we have the following relationships.

$$
\begin{aligned}
(\mathring{\beta} - \tilde{\beta}) &\xrightarrow{p} \frac{\sigma_{1X}}{\sigma_{XX}} - \Pi \left(\frac{\sigma_{XX} - \tau_x}{\sigma_{XX}} \right) \\
\left(\tilde{R} - \mathring{R} \right) \hat{\sigma}_{yy} &\xrightarrow{p} \sigma_{11} + \Pi^2(\tau_x) - \frac{1}{\sigma_{XX}} \left(\sigma_{1X} + \Pi(\tau_x) \right)^2 \\
(R_{max} - \tilde{R})\hat{\sigma}_{yy} &\xrightarrow{p} \Pi \left(\frac{\sigma_{11}\tau_x}{\sigma_{1X}} - \Pi\tau_x \right)
\end{aligned}
$$

These define a system of three equations in three unknowns (σ_{11},σ_{1X} and Π). Solving recursively leaves us with Π as the root of the equation $f(\nu)$ given above. This is a cubic with all real coefficients so it has either one or three real roots. If it has a single real root, that is the solution. If it has multiple real roots, one of the three will be the solution. □

Proposition 2 parallels Theorem 2 in the Altonji, Edler and Taber (2002) working paper; they consider the case where X is binary.

The corollary below develops the case of $\delta = 1$.

Corollary 1. *Define*

$$
\begin{aligned}
\nu_1 &= \frac{-(\Theta) - \sqrt{(\Theta)^2 + 4\left((R_{max} - \tilde{R})\sigma_{yy} \right) \left(\mathring{\beta} - \tilde{\beta} \right)^2 \sigma_{XX}^2 \tau_x}}{-2\tau_X \left(\mathring{\beta} - \tilde{\beta} \right) \sigma_{XX}} \\
\nu_2 &= \frac{-(\Theta) + \sqrt{(\Theta)^2 + 4\left((R_{max} - \tilde{R})\sigma_{yy} \right) \left(\mathring{\beta} - \tilde{\beta} \right)^2 \sigma_{XX}^2 \tau_x}}{-2\tau_X \left(\mathring{\beta} - \tilde{\beta} \right) \sigma_{XX}}
\end{aligned}
$$

where $\Theta = \left(\left((R_{max} - \tilde{R})\sigma_{yy} \right) (\sigma_{XX} - \tau_x) - \left(\left(\tilde{R} - \mathring{R} \right) \sigma_{yy} \right) \tau_x - \sigma_{XX}\tau_x \left(\mathring{\beta} - \tilde{\beta} \right)^2 \right)$. *Define a set* $\beta^* = \{\tilde{\beta} - \nu_1, \tilde{\beta} - \nu_2\}$. *One element of the set* β^* *converges in probability to* β.

Proof. This follows immediately from Proposition 2, with $\delta = 1$. See Appendix A.2. □

In either case – regardless of whether $\delta = 1$ – this problem may have multiple solutions. Only one element of the set will converge in probability to the true β. I discuss solution selection below.

Proposition 3 shows a result related to δ. In particular, I solve for the value of δ to match a particular treatment effect. This will be central to implementation since it allows us to ask how large the relative selection on observables and unobservables would need to be to produce a treatment effect of zero.

Proposition 3. *Define some value* $\hat{\beta}$. *Define* $\hat{\delta}$ *as the coefficient of proportionality for which* $\beta = \hat{\beta}$. *Define:*

$$
\delta^* = \frac{\left(\tilde{\beta} - \hat{\beta} \right) \left(\tilde{R} - \mathring{R} \right) \hat{\sigma}_{yy} \hat{\tau_x} + \left(\tilde{\beta} - \hat{\beta} \right) \sigma_{XX} \hat{\tau}_x (\mathring{\beta} - \tilde{\beta})^2 + 2 \left(\left(\tilde{\beta} - \hat{\beta} \right) \right)^2 \left(\hat{\tau}_x (\mathring{\beta} - \tilde{\beta}) \sigma_{XX} \right) + \left(\left(\tilde{\beta} - \hat{\beta} \right) \right)^3 \left((\hat{\tau}_x \sigma_{XX} - \hat{\tau}_x^2) \right)}{\left((R_{max} - \tilde{R}) \hat{\sigma}_{yy} (\mathring{\beta} - \tilde{\beta}) \sigma_{XX} + \left(\tilde{\beta} - \hat{\beta} \right) (R_{max} - \tilde{R}) \hat{\sigma}_{yy} (\sigma_{XX} - \hat{\tau}_x) + \left(\left(\tilde{\beta} - \hat{\beta} \right) \right)^2 \left(\hat{\tau}_x (\mathring{\beta} - \tilde{\beta}) \sigma_{XX} \right) + \left(\left(\tilde{\beta} - \hat{\beta} \right) \right)^3 \left((\hat{\tau}_x \sigma_{XX} - \hat{\tau}_x^2) \right) \right)}
$$

Under this definition, $\delta^* \xrightarrow{p} \hat{\delta}$.

Proof. The proof follows from setting $\Pi = \tilde{\beta} - \hat{\beta}$, substituting into the $f(\nu)$ function and solving for δ. $\square$

Proposition 3 shows there is a single value of δ to match any targeted treatment effect – for example, a single value of δ will match a treatment effect of zero.

3.2.1 Interpretation of ϵ

In discussing empirical applications, it will be crucial to take a stand on the value of R_{max} which is influenced by ϵ. Conceptually, ϵ represents an error which is uncorrelated with X, W_1 or W_2. One interpretation of ϵ is that it captures the degree of measurement error in the outcome. Another interpretation is that ϵ captures the influence of anything which is determined after X, W_1 and W_2 are determined. Both of these interpretations may be useful in choosing a value for R_{max} in a particular context.

3.2.2 Solution Selection

This estimator may deliver multiple solutions for β. One of these will be the true β under the proportional selection relationship With an added assumption we can typically eliminate at least one solution and, in the case where $\delta = 1$, always produce a single solution.

Define $\hat{W}_1$ as the predicted index of controls from a regression of Y on X and the observed controls. This index uses the coefficients on controls estimated in the regression, which are not the true coefficients. Define the index using the true coefficients as W_1.

Assumption 1: $Sign(Cov(X, \hat{W}_1)) = Sign(Cov(X, W_1))$.

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 1, if $\delta = 1$ there is a unique solution.

In the case where $\delta \neq 1$ there may be multiple solutions, one closer to the controlled treatment effect and one further. The natural heuristic procedure of selecting the treatment effect closest to the controlled coefficient will be appropriate if one is willing to assume the bias is fairly small.

I argue below that in empirical settings a value of $\delta = 1$ is a good bounding value; this is consistent with arguments in Altonji, Elder and Taber (2005). For the purposes of implementation, therefore, it may be appropriate to consider either (a) calculating the bias-adjusted effect under the assumption of $\delta = 1$, with Assumption 1 active or (b) calculating the value of δ such that $\beta = 0$. Either of these will provide a unique solution.

3.2.3 Additional Controls

It is useful to consider a simple extension in which there is an additional observed set of controls which do not share covariance properties with the unobservables but do correlate with X.

Formally, consider the case where the full model is

$$Y = \beta X + \Gamma \mathbf{w^o} + W_2 + m + \epsilon \tag{3}$$

where m is orthogonal to $\mathbf{w^o}$, W_2 and ϵ and the assumptions about orthogonality with ϵ are as above. Assume that the covariance between m and X is unrelated to the covariance between X and $\mathbf{w^o}$ and W_2. It is straightforward to observe in this case that if we simply residualize all other variables with respect to m we return to the setup above and the results go through as stated there. In practice, this means that the controls m are included in both controlled and "uncontrolled" regressions, and X is residualized with respect to m when generating σ_{XX} and $\tau_{x.}$

3.2.4 Inference

Standard errors around β^* could be generated using a bootstrap approach. Such an approach depends on the estimator displaying asymptotic normality. Here, I show evidence for this using simulation. I simulate data from two populations with varying data generating processes.[9] The populations are of size 1,000,000 and I run 1,000 Monte Carlo simulations of the estimator, drawing 10,000 observations each time.

The distributions of estimated β^* in the two cases are shown in Figure 2. A normal distribution is overlaid. The distributions appear normal and a Shapiro-Wilk test does not reject normality in either case. This suggests that a bootstrap approach may be an acceptable way to generate standard errors if that is of interest.

3.3 Relation between Simple and General Estimator

In many cases the simple estimator derived in Section 3.1 may provide an approximation to the general estimator in Section 3.2. Recognizing this – when it is the case – is of value because the simple estimator is intuitive and straightforward to estimate. Conversely, recognizing when this is not the case is useful for identifying scenarios in which coefficient stability may be misleading even in the presence of sizable R-squared movements. I discuss the relation between the estimator, including some simulation evidence, below.

[9] The inputs are described in the figure notes.

Conditions for Simple Estimator to Provide an Approximation

Differences exist between the simpler estimator derived in Section 3.1 and the general estimator even if we consider both under the assumption of equal selection. The reason for this is straightforward. As I note at the end of Section 3.1, in the case of the model with the single observable the intermediate regression recovers the same coefficient and R-squared that one would observe if we could observe and control directly for the index $W_{1.}$ In the case with multiple observables, the intermediate regression does not produce the same values. Therefore, the result is not exact. Consider the two regressions below:

$$
\begin{aligned}
Y &= \tilde{\beta}X + \Gamma W_1 + \tilde{\epsilon} && (4)\\
Y &= \hat{\beta}X + \mathbf{\Psi}\omega^{\mathbf{o}} + \hat{\epsilon} && (5)
\end{aligned}
$$

where W_1 is an index of the elements of $\omega^{\mathbf{o}}$multiplied by their true coefficients. The simple estimator recovers the bias from (4). In order for it to also recover the bias in (5) it must be the case that $\hat{\beta} = \tilde{\beta}$, which will not generically be true.

In cases where $\hat{\beta} \approx \tilde{\beta}$, the simple estimator developed in Section 3.1 will provide an approximation to the general estimator.

Consider a regression of X on $\omega^{\mathbf{o}}$ and denote the coefficients from this regression μ_i. Recall the coefficients on these controls in the regression of Y on X and $\omega^{\mathbf{o}}$ are $\psi_{i.}$

Proposition 4. *If* $\frac{\psi_i}{\psi_j} = \frac{\mu_i}{\mu_j} \forall i, j$ *then the simple estimator is a consistent estimator in the general case given equal selection.*

Proof. Referring to equations (5) and (6) above, Note that $\tilde{\beta} \xrightarrow{p} \beta + \frac{\delta\sigma_{22}\sigma_{1X}}{\sigma_{11}(Var(\tilde{X}))}$ where $\tilde{X}$ is the residual from a regression of X on W_1, and $\hat{\beta} \xrightarrow{p} \beta + \frac{\delta\sigma_{22}\sigma_{1X}}{\sigma_{11}(Var(\hat{X}))}$ where $\hat{X}$ is the residual from a regression of X on $\mathbf{w^o}$. Using the definitions of $plim(\tilde{\beta})$ and $plim(\check{\beta})$ it is straightforward to observe that the result requires $Var(\tilde{X}) = Var(\check{X})$. By the definition of variance, $Var(\tilde{X}) = 1 - \frac{Cov(\Psi\omega^{\mathbf{o}},\mathbf{X})^2}{Var(\Psi\omega^{\mathbf{o}})}$ and $Var(\hat{X}) = 1 - \sum_{i=1}^{j} \mu_i Cov(\omega_i^o, X)$. Algebraic manipulation yields the result that $Var(\tilde{X}) = Var(\check{X})$ if and only if $\frac{\psi_i}{\psi_j} = \frac{\mu_i}{\mu_j}\ \forall i, j$. □

The intuition behind this condition is straightforward: the relative contributions of each variable to X must be the same as their contribution to Y. This will virtually never be absolutely true except in very pathological cases, but in many practical cases the deviation in the estimators is fairly minor. If the effect of the treatment is fairly small, the simple approximation will work well even if this condition is not satisfied. I demonstrate this more concretely in simulation evidence below.

Estimator Deviation: Coefficient Stability in Presence of Large Bias

In cases where the simple estimator does not provide a good approximation, it is key to realize that it is possible for coefficients to appear stable in the presence of large bias even if there is some substantial change in R-squared.

To see this, assume $\delta = 1$ and consider the conditions under which the uncontrolled coefficient $\mathring{\beta}$ is exactly equal to the controlled coefficient $\tilde{\beta}$. Using the notation above, this occurs if and only if $\frac{\sigma_{1X}}{\sigma_{XX}} + \frac{\sigma_{1X}\sigma_{22}}{\sigma_{11}\tau_x}\frac{\tau_x}{\sigma_{XX}} = \frac{\sigma_{1X}\sigma_{22}}{\sigma_{11}\tau_x}$ One condition which will cause this to hold is if $\sigma_{1X} = 0$. The formula for the bias is $\frac{\sigma_{1X}\sigma_{22}}{\sigma_{11}\tau_x}$ so if $\sigma_{1X} = 0$, then there is no bias and $\beta = \tilde{\beta}$.

However, this condition will also hold if $\sigma_{11} = \frac{\sigma_{XX}-\tau_x}{\tau_x}\sigma_{22}$. Under this assumption, the movement in R-squared is $(\sigma_{XX} - \tau_x)\left(\frac{\sigma_{22}}{\tau_x} - \left(\frac{\sigma_{1X}}{(\sigma_{XX}-\tau_x)}\right)^2\right)$ which will be non-zero as long as $\sigma_{XX} > \tau_x$ and $\sigma_{22} > 0$. In this way, the coefficient movement is zero and the R-squared movement is positive, which would appear to suggest limited (or zero) bias. However, the bias in this case is actually $\frac{\sigma_{1X}}{\sigma_{XX}-\tau_x}$ which is non-zero.

Simulated Data

To give a sense of the underlying parameters which would produce these patterns, I present some simulation evidence.

I simulate data from the following model

$$Y = 1 + 200X + \gamma_1 w_1^o + \gamma_2 w_2^o + W_2$$

under varying assumptions about (1) the γ_1 and γ_2 values; (2) the covariance between w_1^o and w_2^o and X; and (3) the variance of W_2. In all cases, I assume $\delta = 1$, $\beta = 200$, $Var(X) = Var(w_1^o) = Var(w_2^o) = 1$ and $R_{max} = 1$. In the results I show first the treatment effects estimated with and without controls. I then show the value of β^* produced by the general estimator using the assumption described above for root selection. In addition, I report the β^* outputted by the simple estimator. I run these deterministically by defining Y within the sample. This means that any differences observed reflect asymptotic bias from the simple approximation.

The results are shown in Table 2. There are five columns, corresponding to five sets of assumptions about the free parameters.

Column (1) adopts the proportionality assumption which is developed above and under which the simple estimator and the general estimator both give the same response. Both estimators produce $\beta^* = 200$, which is the true treatment effect. Column (2) moves slightly away from proportionality, and shows that in doing so the simple estimator is no longer exact. It is extremely close. Column (3) considers a case where the proportionately assumption is seriously violated, and shows that while that change increases the asymptotic bias from the simple estimator, it remains small. In Column (4) I return to the case in Column (2) but increase the variance of W_2 dramatically. This corresponds to a case where the unobservables are relatively more important in explaining Y but the observables relate to Y and X in a similar way. This assumption increases the asymptotic bias in the simple estimator, but the error is still very small.

Finally, in Column (5) I take the assumptions about covariances from Column (3) but assume the variance of W_2 is much larger. This corresponds to a case where the unobservables explain a lot of Y relative to the observables *and* the relationship between w_1^o, X and Y is very different from the relationship between w_2^o, X and Y. In this case, the general estimator estimates the correct β but the simple estimator deviates significantly and, in fact, is the wrong sign. Column (6) demonstrates that the deviation between estimators is erased if the covariance between the controls and X is smaller, even if the proportionality is violated in the same way and the unobservables are important.

The data in Column (5) gives a sense of the assumptions which underlie large deviations between the simple and the general estimator. This will occur in cases where (1) the covariance between the observed controls and treatment is high; (2) the proportionality assumption outlined above is seriously violated and (3) the unobservables are important relative to the observables. A key thing to note is that in the example in Column (5) the coefficient looks fairly stable and there is a sizable move in R-squared. It is only when the full bias adjustment is performed that the bias is revealed.

The theory and simulation evidence make clear the theoretical need to take into account the formal bias calculation. When I turn to empirical work in Sections 5 and 6 I will discuss to what extent these conditions operate in data. I will find that although the simple estimator is generally a good fit there are settings in which the deviation is large.

4 Implementation: Bounding and Robustness Statements

In empirical work in economics, discussions of coefficient stability are typically used in establishing robustness. The estimator above suggests two related ways that such robustness statements might be made. I detail these below.

Statements about δ

One approach to robustness is to assume a value for R_{max} and calculate the value of δ for which $\beta = 0$. This can be interpreted as the degree of selection on unobservables relative to observables which would be necessary to explain away the result (under the full model hypothesized). A value of $\delta = 2$, for example, would suggest that the unobservables would need to be twice as important as the observables to produce a treatment effect of zero.

This approach is akin to the robustness statements suggested by Altonji, Elder and Taber (2005).[10] They suggest that a value of $\delta = 1$ may be a heuristic cutoff. A value of $\delta = 1$ suggests the observables are at least as important as the unobservables. One reason to favor this is that researchers typically focus their data collection efforts (or their choice of regression controls) on the controls they believe *ex ante* are the most important (Angrist and Pischke, 2010). A second is that W_2 is residualized with respect to $\omega^{\mathbf{o}}$ so, conceptually, we want to think of the omitted variables as having been stripped of the portion related to the included ones.

Performing the robustness check in this form requires that researchers make an assumption about R_{max}. One option is to have $R_{max} = 1$, but in many cases this may overstate the total explanatory power of the possible variables. In general, this will be application-specific.

Bounding Statements about β

A second approach to robustness is to use some bounding assumptions on R_{max} *and* δ to develop a set of bounds for β. Such bounds could then be compared to, for example, a value of zero or some other boundary of interest.

I consider this with language similar to partial identification (Tamer, 2010; Manski, 2003). Consider the estimator $\beta^*(R_{max}, \delta)$ defined as above. Without any additional assumptions, I note that R_{max} is bounded between $\tilde{R}$ (the controlled regression R-squared) and 1. I assume that the proportional selection is positive: that is, that the covariance between X and the observables is the same direction as the correlation between X and the unobservables. This bounds the value of δ below at 0 and it is bounded above at some arbitrary upper bound $\overline{\delta}$.

We can then define some bounds for β. One side of the bound is $\tilde{\beta}$, the value of β delivered when $R_{max} = \tilde{R}$ or $\delta = 0$. The other bound is $\beta^*(1, \overline{\delta})$. Without more assumptions, this is either positive or negative infinity, since $\overline{\delta}$ is unbounded. The insight of partial identification is that it may be possible to use additional intuition from the problem to further bound both R_{max} and δ values.

Consider first the issue of bounding δ. I argue that for many problems $\delta = 1$ is an appropriate bound, for the reasons discussed above. Ultimately, this is an empirical issue, and I will discuss at least some evidence for this bound in Section 5.

In the case of R_{max} it may be possible to generate a bound smaller than 1 by, for example, considering measurement error in Y or evaluating variation in Y which cannot be related to X because it results from choices made after X is determined. Define an assumed upper bound on R_{max} as $\overline{R_{max}}$, with $\overline{R_{max}} \leq 1$.

With these two bounding assumptions I can define a bounding "set" as: $\Delta_s = [\tilde{\beta}, \beta^*(\overline{R_{max}}, 1)]$.

Empirically, the question of interest in considering Δ_s is whether the conclusions based on the full set

[10] The calculation will be different since their test produces a value of δ under the null that $\beta = 0$, whereas the calculation here is correct for the true β.

are similar to what we would draw based on observing the controlled coefficient $\tilde{\beta}$. If inclusion of controls moves the coefficient toward zero, one natural question is whether the set includes zero. Regardless of the direction of movement one could ask whether the bounds of the set are outside the confidence interval on $\tilde{\beta}$ – this effectively asks whether the magnitude conclusions based on the controlled coefficient are robust.

This suggested robustness leaves open the question of what is a reasonable $\overline{R_{max}}$ to assume in describing the identified set. I discuss this in two specific empirical contexts in Section 5 and in more detail in the context of the economics literature in Section 6.

Stata Code

Either of these calculations can be performed using STATA code which accompanies this paper. The command is **psacalc.**

5 Empirical Validation

The results above provide a way to recover an estimate of unbiased treatment effects under the assumption that selection on observables and unobservables is proportional. However, the theoretical discussion does not provide any insight as to how this is likely to perform in empirical settings.

In this section I explore this issue using two approaches. In the first subsection, I approach estimator validation by asking how this adjustment performs in constructed data where, by definition, we know the treatment effect. I construct the data with a full set of controls and then explore coefficient bias when various sets of controls are excluded. This allows for a test of whether the proportional selection relationship would lead to better inference in this setting, and allows for direct estimation of values of δ. The latter is helpful in evaluating the empirical validity of the bounding assumption suggested above. I perform this exercise in the familiar setting of wage returns to education.

In the second subsection I use observational data on the relationship between maternal pregnancy and early life behaviors and child outcomes. I compute possibly biased treatment effects, perform the adjustment, and compare the resulting conclusions to external evidence on causal impacts. I ask whether the adjusted coefficients generate more accurate conclusions than the simple controlled estimates.

5.1 Constructed Data: Returns to Education

In this section I consider validation of the estimator in real data which is constructed such that we know the treatment effect. I use the canonical example of estimating wage returns to education.

Estimation of this relationship starts with standard Mincer regressions of wages on education, experience and experience-squared. One central confound is family background: people whose mothers have more education, for example, are more likely to be highly educated but also have higher wages for other reasons.[11] Using data from the NLSY, I construct a dataset in which I define the "true" return to education as the impact of education controlling for a full set of family background characteristics. I then consider the bias – both in simple controlled regressions and after this adjustment is performed – in hypothetical cases in which I do not observe the full set of controls. This exercise will allow me to see how the adjustment performs, to compare the performance of the simple and the general estimator and to estimate values of δ and ask how they compare to the bounds suggested in Section 4.

5.1.1 Data and Empirical Strategy

I use data from the NLSY-79 cohort. I am concerned with the impact of years of education on log wages, and I begin by considering the standard Mincer regression of log wages on educational attainment. I use the higher of the two educational levels recorded in 1981 and 1986 and the higher of the two wage values recorded in 1996 and 1998. Experience and experience-squared are calculated in the typical way (experience = age – education years – 6). I also control for individual sex.

My concern is with confounding by demographics and family background. I capture this with eight variables: region of residence, race, marital status, mother's education, father's education, mother's occupation, father's occupation and number of siblings. All variables are controlled for fully flexibly, with dummies. Summary statistics for these data appear in Appendix B.

I construct a dataset by regressing log wages on education, experience, sex and the full set of family background data. I generate fitted values, and then take these as the "true" effects in the model – that is, the effect on education we see in this regression is the unbiased treatment effect in the constructed data.[12]

The regression of this fitted value on the full set of controls has an R-squared of 1 by construction. In practice, however, wages are not fully predicted by family background or individual characteristics. I therefore add an orthogonal error term to this fitted value. To generate a magnitude for this term I regress the log wage measure used here on log wages in 1992 or 1994 (again, I take the higher of the two). This regression has an R-squared of 0.45. I argue that family background, education, etc, should not explain more of the outcome than the previous year's wages, since these variables all contribute to that wage. I therefore add an orthogonal error term to the fitted value such that the ultimate regression R-squared is about 0.45.

It is important to note that the addition of this error term is done largely for realism; it will be instructive to explore errors that may be introduced by incorrectly assuming that $R_{max} = 1$. However, the

[11] A second obvious issue is the confound with ability. It would be possible to do an exercise similar to this one with that confound. Since the exercise here is not about finding the causal effect of education on wages, but is simply about exploring this adjustment, there is no loss to ignoring the issue of ability.

[12] Clearly, this is not to suggest that this is the causal impact of education on wages. I mean only to assume that this is the true effect in the constructed data, against which I will evaluate estimates which exclude some of the controls used in constructing the effect.

calculations of δ is not sensitive to this addition.

Given this constructed dataset, the empirical exercise is straightforward. I iterate through excluding all sets of controls (up to 6 of the 8). In each case I: (1) calculate the δ implied by the included and excluded control set; (2) calculate β^* with this δ and the true R_{max}; (3) calculate whether the set bounded by $\tilde{\beta}$ and $\beta^*(R_{max}, 1)$ contains the true effect; and (4) calculate β^* with the simple estimator to evaluate the approximation.

5.1.2 Results

Figure 3a shows the distributions of the true β and the estimated $\tilde{\beta}$ and the values of β^*. The true effect in the constructed data is 0.087, with a standard error of 0.003. The β^* values cluster at the true effect value. This is a simple numerical check of the procedure in realistic data: if we know the true R_{max} and the true δ the adjustment works as it should. Not surprisingly, the estimates of $\tilde{\beta}$ are shifted substantially to the right from the true β. Controlled estimates are systematically biased to estimate excess returns to education.

Figure 3b shows the values of δ calculated in this exercise. This value is not mechanical: nothing in the setup constrains any particular value of δ. In the figure, I show the full distribution of δ and the [0,1] bounds that I suggest would be appropriate in many settings.

The average δ is 0.545 and 86% of values fall within the [0,1] range. Only 2 (of 211) values are negative. The cases with values of $\delta > 1$ are instructive. These are combinations of controls where the index of the omitted variables are more important in explaining education than the included ones. Of the 28 cases with $\delta > 1$, 92% of them excluded either maternal or paternal education. This makes clear that these variables are among the most important confounds; this should not be surprising and, indeed, it seems likely that researchers would think to include these first, before considering data on (for example) parental occupation or number of siblings. Put differently, if we consider control set selection not at random as I do here but with the idea that the most important controls are selected first, it is likely that the $[0, 1]$ bound would fit in an even larger share of cases. The fact that the average δ is less than 1 supports the idea of 1 as a bound on δ, rather than as an average value.

I can comment on the bounding logic described in Section 4. Given the δ values, it is straightforward to observe that if we calculate the set $[\tilde{\beta}, \beta^{*\prime}(R_{max}, 1)]$, in 89% of cases this will include the true value. This is an improvement over the simple controlled regression. The naive estimate with controls captures the true value of β only 62% of the time.

As discussed above, it is useful to evaluate how much worse the performance of the adjustment would be if we used the simple approximation rather than the general estimator. Figure 3c replicates 3a but using the simple estimator rather than the general estimator. The figure is extremely similar, suggesting that in this setting the error from using the simple estimator in this case would be small.

As a final point, it is worth saying that if I was to use a value of $R_{max} = 1$ to do these calculations the adjustment would be too large and the effects therefore biased downward. The errors in this case be extremely large.

5.2 Observational Data: Maternal Behavior and Child Outcomes

A second approach to validation is to take a setting in which we have some possibly biased observational relationships and we think we have a sense of the causal effect from external sources. Given this, the question is whether this approach can separate causal from non-causal associations.[13]

In this section I undertake this type of validation exercise in the context of the link between maternal behaviors, infant birth weight and child IQ. These relationships are of some interest in economics, and of wider interest in public health and public policy circles. A literature in economics demonstrates that health shocks while children are in the womb can influence early outcomes and later cognitive skills (e.g. Almond and Currie, 2011; Almond and Mazumder, 2011). A second literature, largely in epidemiology and public health, suggests that even much smaller variations in behavior – occasional drinking during pregnancy, not breastfeeding – could impact child IQ and birth weight. These latter studies, in particular, are subject to significant omitted variable concerns, largely associated with omitted socioeconomic status and family background. I consider five relationships in all: the relationship between child IQ and breastfeeding, drinking during pregnancy, and low birth weight/prematurity and the relationship between birth weight (as the outcome) and maternal drinking and smoking in pregnancy.

5.2.1 Data

I use NLSY data, this time from the Children and Young Adult sample, which has information on the children of NLSY participants. I measure IQ with PIAT test scores for children 4 to 8 and birth weight with birth weight in grams as reported by the mother. In the latter analysis I include all children. In all cases I control for child sex and, with IQ, for their age. These are not considered as part of the confounding set.

The IQ treatments are: months of breastfeeding, any drinking of alcohol in pregnancy and an indicator for being low birth weight and premature (<2500 grams and <37 weeks of gestation). The birth weight treatments are maternal smoking and drinking intensity during pregnancy. I measure family background, the confounding category, with child race, maternal age, maternal education, maternal income and maternal marital status. Summary statistics for these data appear in Appendix B.

[13] Altonji, Elder and Taber (2008) do a version of this for the relationship between survival and catheterization.

5.2.2 Empirical Strategy

I run regressions with and without the controls to extract $\mathring{\beta}$, $\mathring{R}$, $\tilde{\beta}$ and $\tilde{R}$. I adopt a bounding value for R_{max} drawn from within sibling correlations (Mazumder, 2011). In theory, R_{max} should reflect how much of the variation in child IQ and birth weight could be explained if we had full controls for family background; I argue this is the thought experiment approximated by the sibling fixed effect R-squared. The figures are 0.61 for IQ and 0.53 for birth weight.

Given this R_{max} bound, I first calculate the set $[\tilde{\beta}, \beta^*(R_{max}, 1)]$. I also find the value of δ which would produce $\beta = 0$ under the assumed R_{max} and compare this to $\delta = 1$. These two analyses effectively contain the same information.

The conclusions from these robustness calculations are compared to the conclusions we would expect to get if we were able to estimate the full model. To ask whether the adjusted coefficient gets it right, we need to know what the correct answer is.

I use two types of evidence. First, I consider external evidence from randomized trials (where available) and meta-analyses. Randomized evidence suggests that breastfeeding is not linked with full-scale IQ (Kramer et al, 2008) and most evidence does not suggest an impact of occasional maternal drinking on child IQ (see, for example: Falgreen-Eriksen et al, 2012; O'Callaghan et al, 2007).[14] In contrast, low birth weight and prematurity do seem to be consistently linked to low IQ (Salt and Redshaw, 2006), a link which also has a biological underpinning (de Kieviet et al, 2012). Occasional maternal drinking is typically not thought to impact birth weight (Henderson, Gray and Brocklehurst, 2007), but there is better evidence that smoking does (e.g. from trials of smoking cessation programs as in Lumley et al, 2009).

Second, I consider the conclusions one would draw from sibling fixed effects regressions in the NLSY data described above, which provides a more "within sample" test of fully controlling for family background. Of course, sibling fixed effects estimates may be subject to their own concerns about causality, so it is perhaps comforting that the conclusions are the same from either source.

5.2.3 Results

Table 3 reports the results: Panel A shows results on IQ, Panel B on birth weight.

The first column shows treatment effects, standard errors and R-squared values without the socioeconomic status controls. Column 2 shows similar values with the full control set. More breastfeeding is associated with higher IQ in these regressions, and low birth weight is associated with lower child IQ. More maternal drinking appears in these data to be associated with *higher* child IQ later, a finding which is likely to be due to selection give limited biological mechanisms. Both samples show smoking and drinking are

[14] Although the question of whether occasional maternal drinking lowers IQ is a controversial issue, as I show below the observational data here actually estimates *positive* impacts of maternal drinking on IQ, and the fact that those effects are not causal is not a subject of much debate.

associated with lower birth weight. All analyses reported here show significant effects with the full set of controls. Interpreting these results in a naive way, one would conclude that each has a significant link with child outcomes.

Column 3 reports whether external evidence, summarized above, suggests a causal impact. As noted, external evidence supports a relationship between low birth weight and IQ and between smoking and low birth weight but the other relationships do not have broad support. Column 4 shows sibling fixed effects regressions, which result in similar conclusions. The only difference is in the impact of low birth weight on child IQ, where the NLSY regression coefficient is significant only at the 11% level.

Column 5 shows the bounding set, using the R_{max} estimates in the top row of each panel and $\delta = 1$. This procedure performs well. The two cases in which the identified set does not include zero are those where the external evidence suggest significant results. Put differently, if one were to use the rule of accepting the effect as causal only if the identified set excluded zero, this would lead to the same conclusions as the external evidence. In all cases the identified set includes the sibling fixed effect estimates. In Column 6, I calculate the values of δ such that $\beta = 0$. I show the effects confirmed in external data are those which have values of $\delta > 1$ required to produce $\beta = 0$.

Finally, Column 7 computes the bounding set using the simple approximation. The conclusions are similar, although in Panel B we do see more divergence in the estimators. Considering the values in Columns 1 and 2, and the R_{max} value used, it is clear why this is: the contribution of the observables to the R-squared is quite small relative to the hypothesized contribution of the unobservables. As is clear in the simulation, this increases the error in the simple estimator.

There are two final points to make about this analysis. First, similar to the wage analysis above, the average value of δ which matches the adjusted effects to the sibling fixed effect values is less than 1 – it is 0.47 – pointing to the value of 1 as a bound. Second, doing these calculations with a value of $R_{max} = 1$ as the bound would lead us to reject all the associations – including the two which are confirmed in outside data.

The results in this section suggest the robustness framework performs well. It also makes clear the importance of doing formal bias calculations. In this latter example, if we based our analysis only on the size (say, in percent terms) of the coefficient movements we would conclude the link between drinking and low birth weight is much more robust than the link between low birth weight and IQ, since the former moves only 10% and the latter 30%. In fact, the low birth weight and IQ link has more external support. This is confirmed by the identified set conclusions, and mechanically it is reflective of the much larger change in R-squared in the low birth weight – IQ relationship.

6 Application to Economics Literature

I turn now to the application of this approach within the economics literature. I undertake two exercises. First, I ask how stability statements in published papers in economics hold up to a version of this adjustment. I illustrate the robustness of results to this adjustment with varying R_{max} values, assuming $\delta = 1$. In addition, I compare the general estimator to the simple approximation and illustrate cases in which the simple approximation is misleading.

In the second sub-section I use evidence from *randomized* data within economics to develop a bound for $R_{max.}$ This provides one approach to completing a robustness argument.

6.1 Coefficients Stability in Non-Randomized Data

The data for this section comes from the published literature in economics. I extract all papers in the *American Economic Review, Quarterly Journal of Economics*, *The Journal of Political Economy* and *Econometrica* from 2008-2010 with at least 20 citations in the ISI Web of Science, and those from 2011-2013 in the same journals with at least 10 citations. From these papers I extract all results where the researcher explores the sensitivity of the result to a control set. The full sample includes 131 results. Computing the estimator requires observing the coefficients and R-squared values from controlled and uncontrolled regressions, *along with* the variance of the outcomes and treatment, and the residual variance of the treatment after adjusting for controls. The latter value in particular is not accessible from typically published results. I use replication files or researcher inquiries to obtain these values; not surprisingly, this was not possible for all the results identified. The final sample includes 89 results.

The full set of citations used appears in Appendix D.

The empirical exercise here is as follows. I extract the relevant inputs from replication files. Note that in cases where controls are included sequentially, I compare the fewest-controls to the most-controls set. For each result, I calculate the bias-adjusted treatment effect with $\delta = 1$ and varying values of R_{max}.

I consider $R_{max} = 1$ as one bound. I also consider a parametrization of R_{max} as a function of $\tilde{R}$: $R_{max} = min\{\Pi\tilde{R}, 1\}$ with varying values of Π. This function allows for some outcomes to have more measurement error or noise than others, and suggests that the degree of variation accounted for by the observables (including the treatment) may be informative as to the degree accounted for by the unobservables. An alternative would be to use $R_{max} = \tilde{R} + \Pi(\tilde{R} - \mathring{R})$, which captures a similar assumption. I work through this version in Appendix C, and show the conclusions are extremely similar.

Having calculated the identified set using these R_{max} values, I consider two standards for robustness. My primary analysis focuses on the subset of results for which the inclusion of controls moves the coefficient towards zero, and simply asks whether the set includes zero. I also consider whether the bounds of the set fall within +/- 2.8 standard errors of the controlled estimate, an analysis which can be done by including results where controls move the coefficient away from zero. This second standard captures a test of whether the magnitude conclusions from the controlled estimate are shared by the adjusted estimate.

The results appear in Figures 4a and 4b. Figure 4a shows the primary robustness with rejection of zero; Figure 4b uses all results and shows the magnitude test. These graphs show the share of relationships which would survive varying values of Π, with $R_{max} = min\{\Pi\tilde{R}, 1\}$. In either case, I find only about 15% to 20% of results would survive $R_{max} = 1$. Within the others, there is a wide distribution of robustness; some share of results would not survive even quite small differences between $\tilde{R}$ and R_{max}.

To quantify this, Panel A of Table 4 shows the share of results which would survive $R_{max} = 1$ and three values of Π. At least a third of studies would not survive $\Pi = 1.25$. Considering the rejection-of-zero robustness criteria, within this set that is not robust to $\Pi = 1.25$, the average study fails at a value of $\Pi = 1.14$ or, in point estimate terms, a predicted increase in R-squared of 0.05 with inclusion of unobserved controls.

One issue in interpreting these results is that the authors of these papers may not be intending these results as a test of omitted variable bias. To address this, I limit to the large subset of papers in which the authors either explicitly comment on the coefficient stability (since remaining omitted variable bias is the only reason that would matter) or explicitly comment on omitted variable bias. Within this subsample, consider the analog of Columns (1) and (3) of Table 4: 21% of these papers would survive $R_{max} = 1$ and 41% would survive $R_{max} = 1.5\tilde{R}$. This is very similar to the overall sample, suggesting it is not the case that the papers which fail by this criteria do so because this is not the intended test.

In nearly all of the analyses discussed here, the authors discuss only coefficient movements. As noted, this is potentially misleading for two reasons. First, it fails to take into account the R-squared movements and, second, it fails to take into account the possible deviation between the simple approximation and the general estimator. It is informative to consider whether these failures actually matter.

First, consider how the conclusions here would differ from those which rely *only* on coefficient movements. To explore this, I choose an $R_{max} = 1.3\tilde{R}$ cutoff and compare the percent reduction (in absolute value) in coefficients for results which do and do not survive this cutoff. I choose this value because it will be the cutoff I identify later in the analysis of randomized data. Figure 6 shows these results. There is virtually no relationship between the coefficient movements and the result survival, illustrating the fact that coefficient movements alone do not provide much insight about these.

Second, I calculate the general estimator and the simple approximation in all cases. I do this assuming, again, that $R_{max} = 1.3\tilde{R}$. In most cases the approximation is fairly close. In 80% of cases the conclusion about robustness is the same. However, in 90% of cases the simple approximation understates the bias, on average by about 40%. The divergence between the general estimator and the simple approximation occurs when two factors converge. First, when the covariance between the treatment and the controls is very large.

Second, when the increase in R-squared with inclusion of controls is small relative to the anticipated change in R-squared with unobservables included.

This insight may be helpful since the first of these, in particular, is straightforward to test in the data. If much of the treatment variance is explained by controls, it may be a signal that coefficient stability is misleading. I will return to this discussion when I discuss some examples in Section 6.3.

6.2 Evidence on Stability Cutoffs from Randomized Data

The evidence above makes clear that, even within a sample of papers which argue for coefficient stability, there is a lot of variation in the robustness of results depending on R_{max}. A natural following question is whether we can suggest any guidance about where one might draw the line – specifically, is there some value of Π (where $R_{max} = \Pi\tilde{R}$) above which we should consider a result robust?

I argue that one place to look for such guidance is in reports from randomized data. Randomized experiments are becoming increasingly common within economics and papers reporting results of these experiments often include regressions with and without controls. Sometimes these are explicitly used to test balance in the experiment, although it is also commonly done to increase precision. Assuming that the data is correctly randomized, if the sample size were infinite, the effects would not be expected to move at all. In practice, with finite data, coefficients can move a bit simply due to very small differences across groups.

When non-randomized papers invoke a coefficient stability heuristic to argue the results they observe are causal, they are (perhaps implicitly) suggesting that the treatment is as good as random. Including controls does not change the coefficient because there is no confounding; this is exactly the argument we know holds in randomized cases. Given this, I argue we can use the stability of randomized data as a guide to how much stability we would expect in non-randomized data *if the treatment were assigned exogenously:* is the coefficient stability within the range the researcher would expect with a randomly assigned treatment?

The approach in this section is to assume effects estimated in randomized data are causal and to therefore assume that they should survive this adjustment procedure.[15] I then ask what value of Π in the R_{max} parametrization would make this true.

The baseline set of papers for this analysis is all randomized papers (lab or field) published in the *American Economic Review, Quarterly Journal of Economics*, *Journal of Political Economy, Econometrica* and the *American Economic Journal – Applied Economics* in the period 2008 through 2013.[16] I extract from these all papers which report sensitivity of a treatment effect to controls. In cases where there are multiple effects reported (i.e. multiple outcomes), I include all effects. I use replication files or researcher requests to

[15] An obvious concern is that, perhaps, these papers are not correctly randomized. This would lead me to a standard which was too lax. I address this in two ways. First, I have focused on papers published in highly ranked journals, increasing the chance that the randomization was of high quality. Second, I will draw guidelines which fit nearly all but not all papers, thus accepting that a small share of randomized papers may suffer from true lack of balance and should not be used to guide this approach.

[16] I include AEJ-Applied because it has published a large number of experimental papers. This journal began publishing in 2009.

extract the estimator inputs where possible. The final sample includes 65 results.

The full set of references is in Appendix D.

I undertake the same analysis as in the non-randomized data: calculate the bias-adjusted treatment effect assuming $\delta = 1$ and varying R_{max} and compare the results to the two standards for robustness.

Figures 5a and 5b show the distributions of sensitivity for the randomized data. A first thing to note is that these results are more robust than the non-randomized results. I have graphed them on the same scale for comparability. About 40% of randomized results would survive a cutoff of $R_{max} = 1$. Nearly all would survive a cutoff of $R_{max} = 1.25\tilde{R}$, much greater than for the non-randomized results. Panel B of Table 4 shows the survival shares for this dataset explicitly under the varying R_{max} cutoffs.

It is not surprising that the randomized results are more robust. The fact that they do not all survive $R_{max} = 1$ is due to the fact that even small changes in coefficients can be blown up with this assumption. I use these data to develop robustness cutoff values. I base these on the value of Π which would allow 90% of results to survive in both the confidence interval and the rejection of zero test. This leads to the bounding values of $\Pi = 1.3$. This value suggests a bound where the unobservables explain somewhat less than the observables (where the latter includes the treatment). This has some intuitive appeal if we think that the observables are chosen with an eye to those which are most important in explaining the outcome.

To argue for a level of stability which would be expected from a randomized treatment, non-randomized effects should show that the set $[\tilde{\beta},\beta^*(min\{1.3\tilde{R},1\},1)]$ excludes zero or, equivalently, that the δwhich produces $\beta = 0$ with $R_{max} = 1.3\tilde{R}$ exceeds 1. Applying this to the non-randomized data above, I find that 51% of results would survive this standard. This standard would be valuable to explore even in cases where the controls cause the coefficient to move away from zero; in that case the question would be whether considering the full set would lead to very different conclusions than the controlled estimate.

6.3 Examples

Before concluding it is informative to consider some examples of when performing this adjustment would impact conclusions. In the first example below I compare several results with similar coefficient movements but very different R-squared movements. In the second I discuss a case in which the simple approximation is far from the general estimator. In all cases I consider these examples using $\delta = 1$ and the bounding value for R_{max} developed based on the randomized data described above.

Importance of R-Squared Movements

To illustrate the dual role of coefficient stability and R-squared movements, I identify three results from the set of papers above which have similar percent coefficient movements with inclusion of controls but vary in their R-squared changes. The three cases are chosen not for their topical similarity but simply for the similar coefficient changes.

The first result is drawn from Lavy (2009), who considers the impact of teacher performance pay on test scores. The result used is the effect on average math scores. The second result is drawn from Brunnermeier and Nagel (2008), who consider whether wealth fluctuations generate changes in portfolio allocations. The result used is the relationship between the change in financial wealth and the share of wealth invested in risky assets. The third result is drawn from Olken and Barron (2009) who explore the relationship between checkpoints on Indonesian roads and bribes paid. The result used is the relationship between the total payments and number of checkpoints.

Table 5 shows in the first columns the coefficient and R-squared values with and without controls. I then report the percentage change in coefficients when controls are added. The three cases all have quite stable coefficients, with coefficient changes in the range of 3% to 5.5%. The final column shows the bias-adjusted β under the proportional selection assumption with $\delta = 1$ and $R_{max} = 1.3\tilde{R}$. The cases differ significantly in the impact of this adjustment on their conclusions. For both Lavy (2009) and Olken and Barron (2009) the bias adjustment makes a fairly small difference. However, for the Brunnermeier and Nagel (2008) result, the difference in coefficients is very large.

This distinction arises due to the varying movements in R-squared values across the three cases. The two cases with small bias adjustments are those where the percentage change in R-squared with inclusion of controls is very big.

Simple Approximation Error

In many cases the simple approximation is close to the general estimator, and taking into account both coefficient movements and R-squared movements in a simple way will provide a good bias estimate. However, this is not always the case and it is informative to consider an example where these deviate significantly and the simple adjustment would be misleading.

Nunn and Wantchekon (2011) analyze the impact of the slave trade on mistrust in Africa. This is a salient example because the authors worry explicitly about unobserved differences across areas, and present a number of arguments to support the interpretation of their results as causal. In contrast to most papers in this literature, they undertake direct calculations based on the theory in Altonji, Elder and Taber (2005). They use coefficient movements in their regressions to calculate the value of δ which would be required to produce $\beta = 0$. They argue the results are robust because all the calculated values of δ are greater than 1. Equivalently, the adjusted treatment effects have the same sign as their controlled effects if $\delta = 1$.

Although it is not made explicit, the calculations they undertake in the paper implicitly assume that

$R_{max} = \tilde{R} + (\tilde{R} - \mathring{R})$.[17] In words, they assume the unobservables explain as much of the outcome as the observables (ignoring that the treatment is part of the observables). In practice, the R-squared values in their regressions do not move much; as an example, in the first row of their Table 4, considering the "Trust Relatives" measure, adding controls increases R-squared from 0.115 to 0.133. Their adjustment assumes that the fully controlled R-squared would be 0.151. The cutoff generated based on the randomized data above would suggest a value of 0.173.

A set of results from their Table 4 are reported in Table 6. The first columns show their estimated effects, and the third column generates the simple approximation to β using their implicit assumption on R_{max} and $\delta = 1$. As stated in their paper, the effects here are of very similar magnitude to the main controlled effects. The fourth column shows the simple approximation of β using the assumption that $R_{max} = 1.3\tilde{R}$ and $\delta = 1$. The adjustments are larger here, but the conclusions are quite similar.

The final column shows the β estimated using the full estimator rather than the approximation (along with the assumptions that $R_{max} = 1.3\tilde{R}$ and $\delta = 1$). The results are very different. In four of five cases the bias-adjusted β is of a different sign than their controlled effects and typically quite large in magnitude. It is worth noting that even if I adopt their conservative assumption on R_{max} two of the five results would show bias-adjusted values above zero (results available from the author).

In looking at the raw data it is clear why this is the case. The controls used explain a large share of the variation in the treatment. This suggests a high covariance between treatment and controls which, in combination with the relatively small R-squared movements, generates a large error in the simple approximation.

7 Conclusion

This paper develops a formal language for discussing robustness of treatment effects, related to the popular heuristic of exploring coefficient sensitivity to controls. I connect this heuristic to the idea of a proportional selection relationship on observed and unobserved variables (Altonji, Elder and Taber, 2005). I describe an implementation strategy for generating bounds on treatment effects and show validation in two empirical contexts. Applying this to the economics literature, and drawing guidelines for expected coefficient sensitivity from randomized results, I develop a full bounding argument.

I suggest a standard for robustness relying on this estimator which could be easily implemented by researchers. A key issue is the need to make an assumption about the share of the outcome variance would be explained by observed and unobserved variables together. I suggest a standard based on the performance of this estimator in randomized data.

[17] They draw this from Bellows and Miguel (2009).

This provides one general approach to developing intuition about R_{max} but it is worth noting that within a given context it may be possible to develop a better intuition. Some examples of this are provided earlier in the paper. In the case of education and wages I develop a value of R_{max}(used for constructing the data) by looking at how much of current year wages are explained by past year wages; the theory is that any ability/motivation/family background confounds are determined prior to the previous years wages (see Section 3.2.1). In the analysis of maternal behavior and child outcomes I use sibling correlations as a benchmark since sibling share the same family background. In two papers following on their original paper (Altonji et al, 2006; Altonji et al, 2008) Altonji and coauthors suggest two methods for adjusting for idiosyncratic variance, an approach parallel to my use of $R_{max.}$

The core insight here is to recognize that coefficient stability on its own is at best uninformative and at worst very misleading. It must be combined with information about R-squared movements to develop an argument. In addition, in the common empirical case in which the researcher is including multiple observed controls, it is possible to observe very stable coefficients and sizable R-squared movements even in the presence of large bias. This strongly argues for approaching these robustness arguments formally.

The robustness approach in this paper addresses concerns related to unobservables which are related to the observables. A key issue which must still be addressed is the appropriate choice of observables (as discussed in Angrist and Pishke, 2010). If there are unobservables which do not share covariance properties with the observed variables then this approach breaks down. Recognizing this issue may help improve the control sets used in empirical work.

References

Almond, Douglas and Bhashkar Mazumder, "Health Capital and the Prenatal Environment: The Effect of Ramadan Observance during Pregnancy," *American Economic Journal: Applied Economics*, October 2011, *3* (4), 56–85.

___ **and Janet Currie**, "Killing Me Softly: The Fetal Origins Hypothesis," *Journal of Economic Perspectives*, Summer 2011, *25* (3), 153–72.

Altonji, Joeseph, Timothy Conley, Todd Elder, and Christopher Taber, "Methods for Using Selection on Observed Variables to Address Selection on Unobserved Variables," 2011. Mimeo, Yale University.

Altonji, Joseph G., Todd E. Elder, and Christopher R. Taber, "An Evaluation of Instrumental Variable Strategies for Estimating the Effects of Catholic Schooling," *Journal of Human Resources*, 2005, *40* (4), 791–821.

___ **,** ___ **, and** ___ , "Selection on Observed and Unobserved Variables: Assessing the Effectiveness of Catholic Schools," *Journal of Political Economy*, 2005, *113* (1), 151–184.

___ **, Todd Elder, and Christopher R. Taber**, "Using Selection on Observed Variables to Assess Bias from Unobservables When Evaluating Swan-Ganz Catheterization," *American Economic Review*, 2008, *98* (2), 345–50.

Angrist, Joshua D. and Jorn-Steffen Pischke, "The Credibility Revolution in Empirical Economics: How Better Research Design Is Taking the Con out of Econometrics," *Journal of Economic Perspectives*, Spring 2010, *24* (2), 3–30.

Bellows, John and Edward Miguel, "War and local collective action in Sierra Leone," *Journal of Public Economics*, December 2009, *93* (11-12), 1144–1157.

Brunnermeier, Markus K. and Stefan Nagel, "Do wealth fluctuations generate time-varying risk aversion? Micro-evidence on individuals' asset allocation," *AMERICAN ECONOMIC REVIEW*, JUN 2008, *98* (3), 713–736.

Chiappori, Pierre-Andrei, Sonia Oreffice, and Climent Quintana-Domeque, "Fatter Attraction: Anthropometric and Socioeconomic Matching on the Marriage Market," *Journal of Political Economy*, 2012, *120* (4), 659 – 695.

de Kieviet, J. F., L. Zoetebier, R. M. van Elburg, R. J. Vermeulen, and J. Oosterlaan, "Brain development of very preterm and very low-birthweight children in childhood and adolescence: a meta-analysis," *Dev Med Child Neurol*, Apr 2012, *54* (4), 313–323.

Dehejia, Rajev and Sadek Wahba, "Causal Effects in Nonexperimental Studies: Reevaluating the Evaluation of Training Programs," *Journal of the American Statistical Association*, 1999, *94* (448), 1053–1062.

Falgreen-Eriksen, H. L., E. L. Mortensen, T. Kilburn, M. Underbjerg, J. Bertrand, H. Stavring, T. Wimberley, J. Grove, and U. S. Kesmodel, "The effects of low to moderate prenatal alcohol exposure in early pregnancy on IQ in 5-year-old children," *BJOG*, Sep 2012, *119* (10), 1191–1200.

Gelbach, Jonah B, "When do covariates matter? And which ones, and how much?," *Journal of Labor Economics*, forthcoming.

Heckman, James and Joseph Hotz, "Choosing Among Alternative Nonexperimental Methods for Estimating the Impact of Social Programs: The Case of Manpower Training," *Journal of the American Statistical Association*, 1989, *84* (408), 862–874.

Henderson, J., R. Gray, and P. Brocklehurst, "Systematic review of effects of low-moderate prenatal alcohol exposure on pregnancy outcome," *BJOG*, Mar 2007, *114* (3), 243–252.

Imbens, Guido, "Sensitivity to Exogeneity Assumptions in Program Evaluation," *American Economic Review*, 2003, *93* (2), 126–132.

Kramer, M. S., F. Aboud, E. Mironova et al., "Breastfeeding and child cognitive development: new evidence from a large randomized trial," *Arch. Gen. Psychiatry*, May 2008, *65* (5), 578–584.

Lacetera, Nicola, Devin G. Pope, and Justin R. Sydnor, "Heuristic Thinking and Limited Attention in the Car Market," *American Economic Review*, August 2012, *102* (5), 2206–36.

LaLonde, Robert J, "Evaluating the Econometric Evaluations of Training Programs with Experimental Data," *American Economic Review*, September 1986, *76* (4), 604–20.

Lavy, Victor, "Performance Pay and Teachers' Effort, Productivity, and Grading Ethics," *AMERICAN ECONOMIC REVIEW*, DEC 2009, *99* (5), 1979–2011.

Leamer, Edward, *Specification Searches: Ad Hoc Inference with Nonexperimental Data*, Wiley, 1978.

Leamer, Edward E, "Let's Take the Con Out of Econometrics," *American Economic Review*, March 1983, *73* (1), 31–43.

Lumley, J., C. Chamberlain, T. Dowswell, S. Oliver, L. Oakley, and L. Watson, "Interventions for promoting smoking cessation during pregnancy," *Cochrane Database Syst Rev*, 2009, (3), CD001055.

Manski, C.F., *Partial Identification of Probability Distributions* Springer Series in Statistics, Springer, 2003.

Mazumder, Bhashkar, "Family and Community Influences on Health and Socioeconomic Status: Sibling Correlations Over the Life Course," *The B.E. Journal of Economic Analysis & Policy*, 2011, *11* (3), 1.

Murphy, Kevin and Robert Topel, "Efficiency Wages Reconsidered: Theory and Evidence," in "Advances in the Theory and Measurement of Unemployment" 1990, pp. 204–240.

Nunn, Nathan and Leonard Wantchekon, "The Slave Trade and the Origins of Mistrust in Africa," *American Economic Review*, 2011, *101* (7), 3221–3252.

O'Callaghan, F. V., M. O'Callaghan, J. M. Najman, G. M. Williams, and W. Bor, "Prenatal alcohol exposure and attention, learning and intellectual ability at 14 years: a prospective longitudinal study," *Early Hum. Dev.*, Feb 2007, *83* (2), 115–123.

Olken, Benjamin A. and Patrick Barron, "The Simple Economics of Extortion: Evidence from Trucking in Aceh," *JOURNAL OF POLITICAL ECONOMY*, JUN 2009, *117* (3), 417–452.

Pearl, Judea, *Causality: Models, Reasoning and Inference*, Cambridge university Press, 2000.

Pischke, Jorn-Steffen and Hannes Schwandt, "Poorly Measured Confounders are Useful on the Left But Not on the Right," September 2013. London School of Economics Working Paper.

Rosenbaum, Paul and Don Rubin, "Assessing Sensitivity to an Unobserved Binary Covariate in an Observational Study with Binary Outcome," *Journal of the Royal Statistical Society, Series B*, 1983, *45* (2), 212–218.

Salt, A. and M. Redshaw, "Neurodevelopmental follow-up after preterm birth: follow up after two years," *Early Hum. Dev.*, Mar 2006, *82* (3), 185–197.

Tamer, Elie, "Partial Identification in Econometrics," *Annual Review of Economics*, 09 2010, *2* (1), 167–195.

Figure 1: **Coefficient Stability and R-Squared Movements**

(a) All Significant Relationships

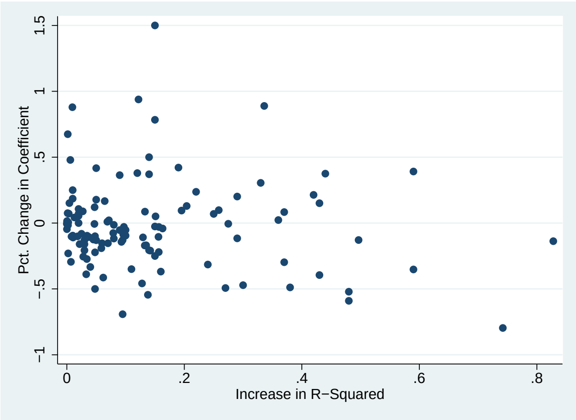

(b) Sample where Controls Lower Magnitude

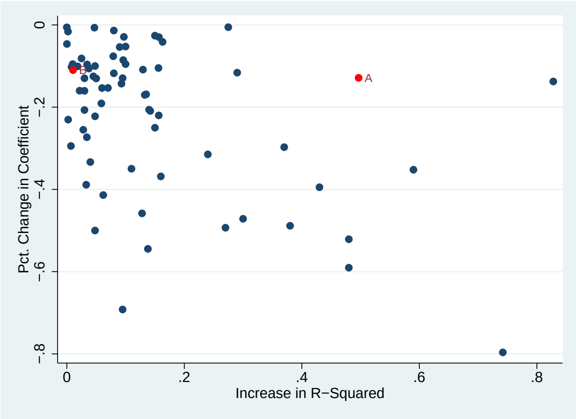

Notes: These figures show the relationship between the percent change in coefficient and the increase in R-squared in sample of highly cited papers from top journals in economics. The sample is discussed in Section 2.

Figure 2: **Distribution of Estimated Treatment Effects**

(a) First Parameter Set

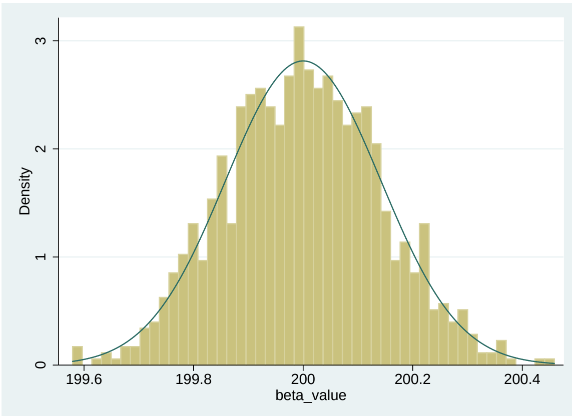

(b) Second Parameter Set

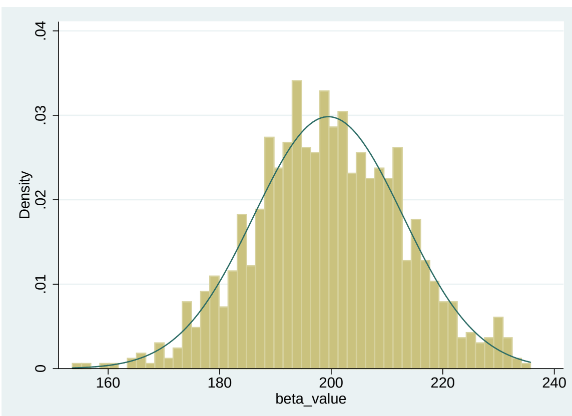

Notes: These figures show the distribution of estimated bias-adjusted treatment effects under two parameter sets. The figures are generated by drawing 1,000 samples of size 10,000 from a population of 1,000,000. The data generating values for the first set are: $\beta = 200, \delta = 1, \gamma_1 = 100, \gamma_2 = 200, Cov(X, w_1^o) = 0.1, Cov(X, w_2^o) = 0.1, Var(W_2) = 20,000, Var(w_1^o) = Var(w_2^o) = 1$. The second set uses the same inputs but with $Var(W_2) = 250,000$. In both cases, I add an *iid* error with mean 100 and standard deviation 1.

Figure 3: **NLSY Wage Data Simulation**

(a) "True", Controlled and Adjusted Beta

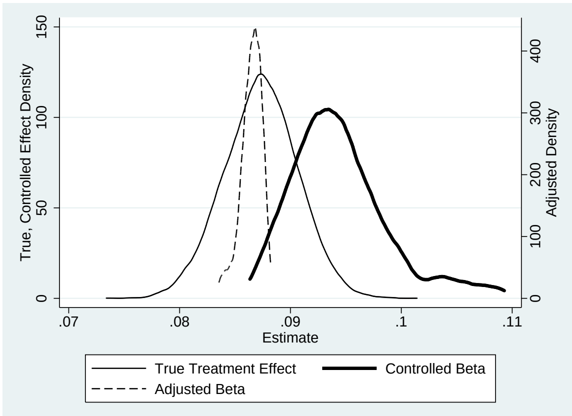

(b) Distribution of Estimated Delta

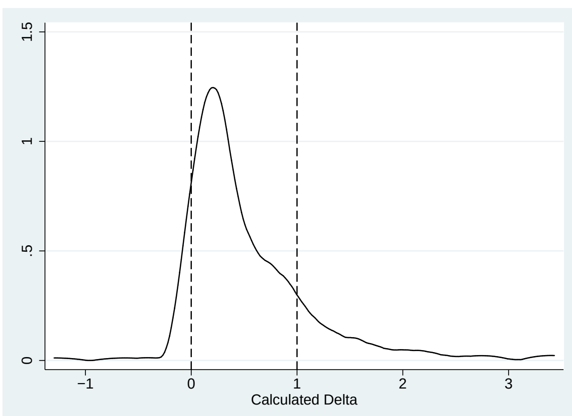

(c) "True", Controlled and Simple Approximation Adjusted Beta

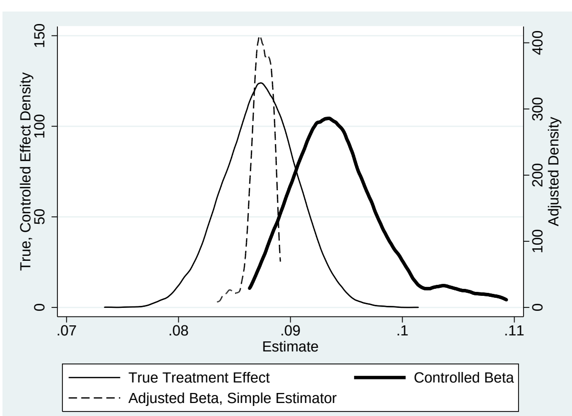

Notes: These figures show results from the validation using the constructed NLSY wage dataset. The analysis is described in Section 5.

Figure 4: **Robustness of Stability Results in Economics Literature**

(a) Rejection of Zero, $R_{max} = \Pi\tilde{R}$.

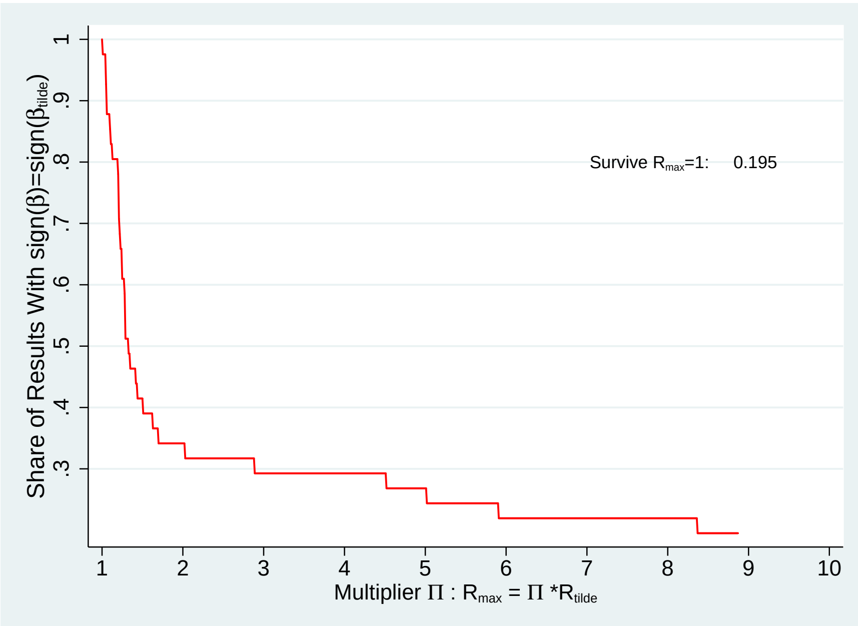

(b) Results within +/- 2.8 SE, $R_{max} = \Pi\tilde{R}$.

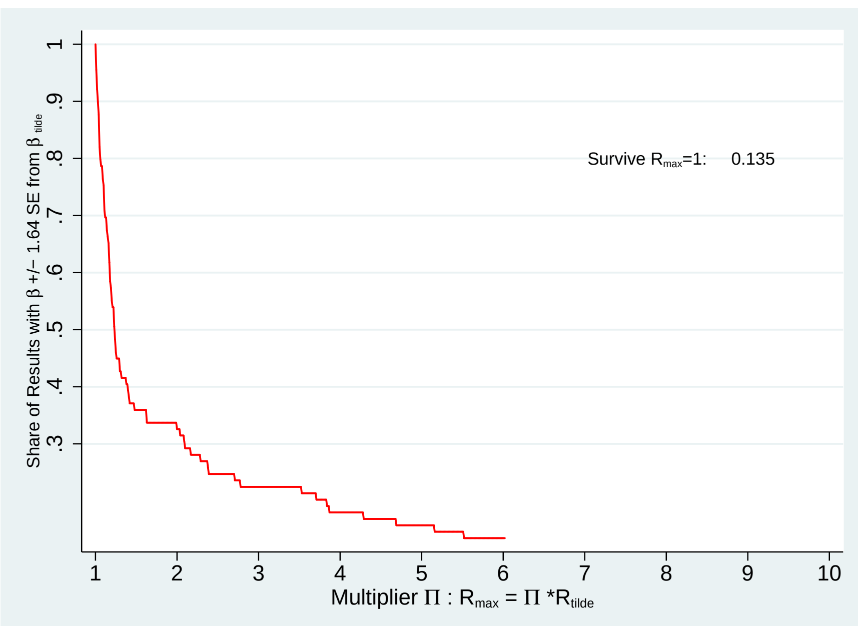

Notes: These graphs show the performance of non-randomized results under the proportional selection adjustment. Each figure graphs the share of results which would survive varying parametrization of R_{max}, in all cases assuming $\tilde{\delta} = 1$. Sub-Figure a indicates the share of results which would survive $R_{max} = \Pi\tilde{R}$ for varying values of Π, with survival in this case meaning the identified set does not include zero. This figure contains only relationships where the effect is significant with controls and adding the controls moves the coefficient toward zero. Sub-Figure b indicates the share of results for which the full identified set would be within 2.8 standard errors of the controlled coefficient. This Sub-Figure includes all relationships.

Figure 5: **Results from Randomized Data**

(a) Rejection of Zero, $R_{max} = \Pi\tilde{R}$.

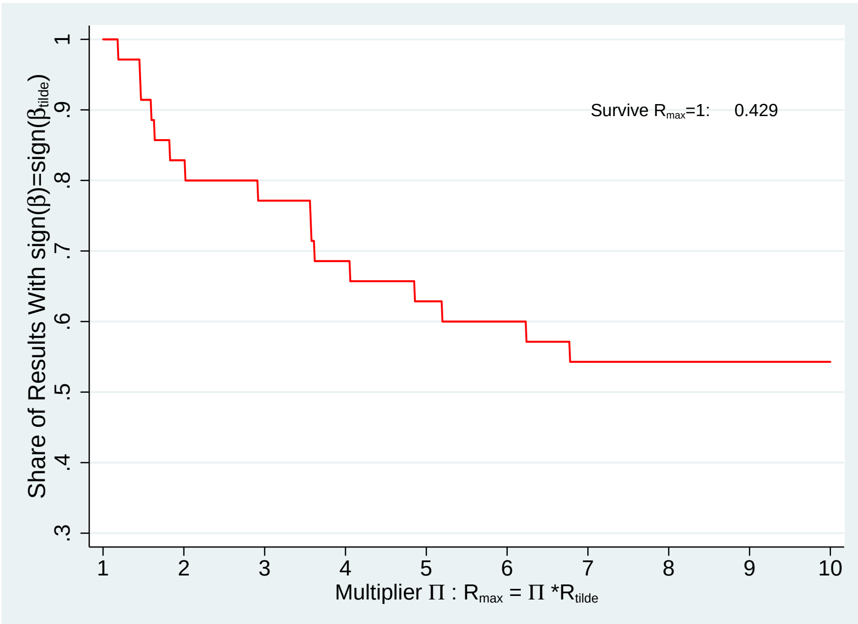

(b) Results within +/- 2.8 SE, $R_{max} = \Pi\tilde{R}$.

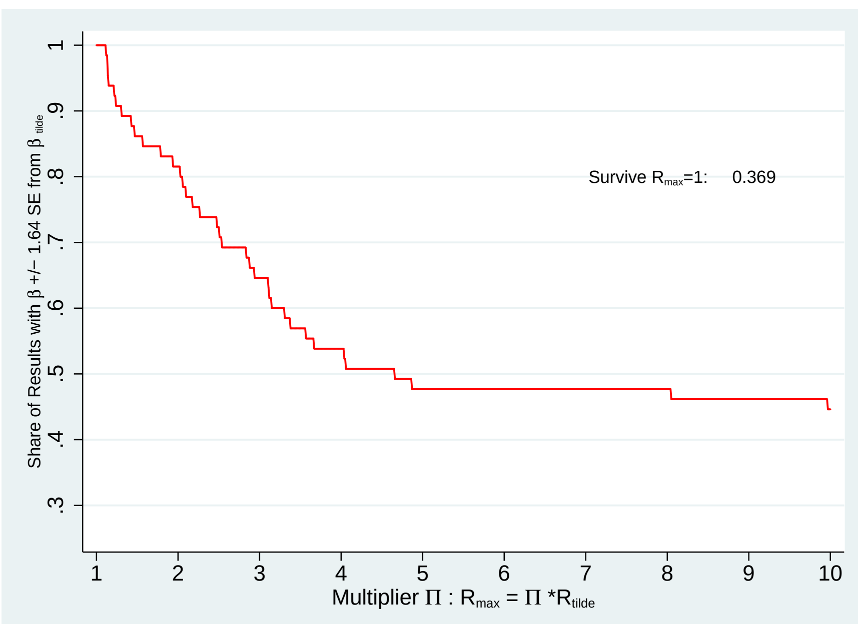

Notes: These graphs show the performance of randomized results under the proportional selection adjustment. Each figure graphs the share of results which would survive varying parametrizations of R_{max}, in all cases assuming $\tilde{\delta} = 1$. Sub-Figure a indicates the share of results which would survive $R_{max} = \Pi\tilde{R}$ for varying values of Π, with survival in this case meaning the identified set does not include zero. This figure contains only relationships where the effect is significant with controls and adding the controls moves the coefficient toward zero. Sub-Figure b indicates the share of results for which the full identified set would be within 2.8 standard errors of the controlled coefficient. This Sub-Figure includes all relationships.

Figure 6: **Relationship between Full Robustness and Coefficient Movement**

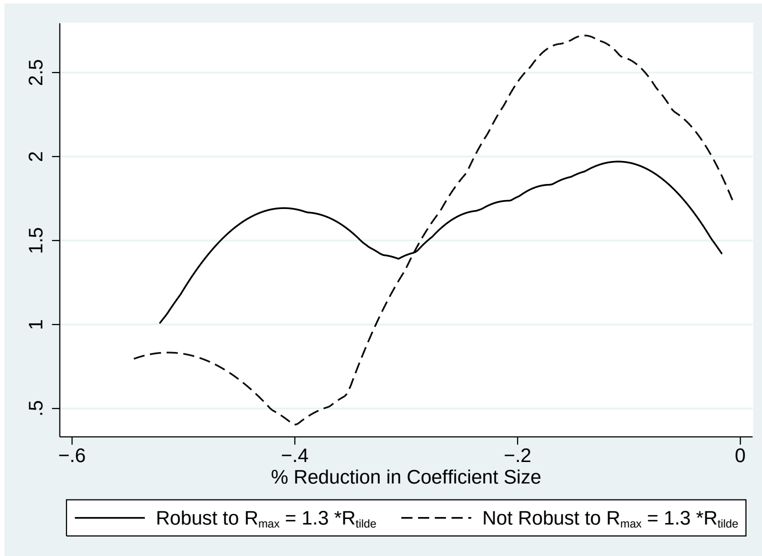

Notes: This graph shows the range of coefficient movements in non-randomized studies. These studies are divided into those which are robust to the proportional selection adjustment with $R_{max} = 1.3\tilde{R}$ (solid line) and those which are not (dotted line). This includes only relationships in which the inclusion of controls moves the coefficient toward zero.

Table 1: **Calibrated Examples**

Panel A: High Versus Low Variance Control			
Quality of Observed Control	Uncontrolled Coefficient. [R^2]	Controlled Coefficient [R^2]	True Effect
High Variance Control Observed	0.202 [.004]	0.002 [.990]	0
Low Variance Control Observed	0.202 [.004]	0.200 [.013]	0

Panel B: Varying Coefficient Movements			
Control Importance in Explaining Y	Uncontrolled Coefficient. [R^2]	Controlled Coefficient [R^2]	True Effect
Very Important	0.20 [.004]	0.195 [.95]	0.194
Not Important	0.20 [.004]	0.195 [.01]	-0.63

Notes: Panel A of this table shows calculations based on simulated data. The true model is $Y = \beta X + W + C$, with $\beta = 0$. The data is constructed so the high variance control is W and $Var(W) = 10$ and the low variance controls is C and $Var(C) = 0.1$. $Var(X) = 1$, $Cov(X, W) = .2$ and $Cov(X, C1) = 0.002$. This last is imposed by the equal selection assumption. Panel B shows an example with identical coefficient movements and varying R-squared movements. The true effect is calculated as the effect under the assumption that the unobserved and unobserved variables together would fully explain Y and the relationship between the observed variable and X is the same as the relationship between the unobserved variables and X.

Table 2: **Simulated Data**

	(1)	(2)	(3)	(4)	(5)	(6)
	$\beta = 200$	$\beta = 200$	$\beta = 200$	$\beta = 200$	$\beta = 200$	$\beta = 200$
	$\gamma_1 = 100$	$\gamma_1 = 100$	$\gamma_1 = -100$	$\gamma_1 = 100$	$\gamma_1 = -100$	$\gamma_1 = -100$
Assumptions:	$\gamma_2 = 200$	$\gamma_2 = 200$	$\gamma_2 = 200$	$\gamma_2 = 200$	$\gamma_2 = 200$	$\gamma_2 = 200$
	$Cov(X, w_1^o) = 0.05$	$Cov(X, w_1^o) = 0.1$	$Cov(X, w_1^o) = -0.1$	$Cov(X, w_1^o) = 0.1$	$Cov(X, w_1^o) = -0.1$	$Cov(X, w_1^o) = -0.01$
	$Cov(X, w_2^o) = 0.1$	$Cov(X, w_2^o) = 0.1$	$Cov(X, w_2^o) = -0.35$	$Cov(X, w_2^o) = 0.1$	$Cov(X, w_2^o) = -0.35$	$Cov(X, w_2^o) = -0.035$
	$Var(W_2) = 200$	$Var(W_2) = 200$	$Var(W_2) = 200$	$Var(W_2) = 250,000$	$Var(W_2) = 250,000$	$Var(W_2) = 250,000$
$\mathring{\beta}[\mathring{R}]$	225.1 [0.50]	230.1 [0.51]	139.7 [0.30]	380.0 [0.35]	-160.0 [0.13]	164.0 [0.08]
$\tilde{\beta}[\tilde{R}]$	200.1 [0.99]	200.1 [0.99]	199.7 [0.99]	353.1 [0.44]	-145.82 [0.25]	169.9 [0.24]
General Estimator β	$\beta^* = 200$	$\beta^* = 200$	$\beta^* = 200$	$\beta^* = 200$	$\beta^* = 200$	$\beta^* = 200$
[Alternative Roots]	[2200]	[1866.6]	[-633.3]	[1866.6]	[-633.3]	[-8133.3
Simple Estimator β	$\beta^* = 200$	$\beta^* = 200.00025$	$\beta^* = 199.98$	$\beta^* = 202.26$	$\beta^* = -59.9$	$\beta = 199.89$

Notes: This table uses the various estimators proposed in Monte Carlo simulations. Results report the asymptotic value of the estimates. All simulations assume δ=1, $R_{max} = 1$ and the variance of w_1^o and w_2^o are both equal to 1. *

Table 3: **Maternal Behavior, Child IQ and Birth Weight**

	(1)	(2)	(3)	(4)	(5)	(6)	(7)
Panel A: Child IQ, Standardized (NLSY) ($R_{max} = .61$)							
Treatment Variable	*Baseline Effect (Std. Error), $[R^2]$*	*Controlled Effect (Std. Error), $[R^2]$*	*Null Reject? (extrnl. evid.)*	*Sibling FE Estimate*	*Identified Set*	δ for $\beta = 0$ *Given R_{max}*	*Identified Set (Simple Approx.)*
Breastfeed (Months)	0.045^{***} (.003) [.045]	0.017^{***}(.002) [.256]	No	-0.007 (.005)	[-0.033,0.017]	0.37	[-0.028,0.017]
Drink in Preg. (Any)	0.176^{***}(.026) [.008]	0.050^{**}(.023) [.249]	No	0.026 (.036)	[-0.146,0.050]	0.26	[-0.138,0.050]
LBW + Preterm	-0.188^{***}(.057) [.004]	-0.125^{***}(.050) [.251]	Yes	-0.111 (.070)	[-0.124,-0.033]†	1.37	[-0.124,-0.033]†
Panel B: Birth Weight in Grams (NLSY) ($R_{max} = .53$)							
Treatment Variable	*Baseline Effect (Std. Error), $[R^2]$*	*Controlled Effect (Std. Error), $[R^2]$*	*Null Reject? (extrnl. evid.)*	*Sibling FE Estimate*	*Identified Set*	$\tilde{\delta}$ for $\beta = 0$ *Given R_{max}*	*Identified Set (Simple Approx.)*
Smoking in Preg	-183.1^{***} (12.9) [.31]	-172.5^{***}(13.3) [.35]	Yes	-94.3^{***}(27.6)	[-172.5,-30.3]†	1.08	[-172.5,-115.7]†
Drink in Preg. (Amt)	-16.7^{***}(5.15) [.30]	-14.1^{***} (5.06) [.34]	No	-1.53 (7.48)	[-14.1,0.49]	0.96	[-14.1,-1.05]†

Notes: This table shows the validation results for the analysis of the impact of maternal behavior on child birth weight and IQ. Baseline effects include only controls for child sex and (1) age dummies in the case of IQ and (2) gestation week in the case of birth weight. Full controls: race, age, education, income, marital status. Sibling fixed effects estimates come from NLSY in all panels. The identified set in Column (5) is bounded below by $\tilde{\beta}$ and above by β^* calculated based on R_{max} given in the top row of each panel and $\tilde{\delta} = 1$.The R_{max} calculation in Column (6) is done under the assumption that $\tilde{\delta} = 1$ * significant at 10% level, ** significant at 5% level, *** significant at 1% level. †identified set excludes zero.

Table 4: **Robustness of Stability Results**

Panel A: Non-Randomized Data, Share of Results which Survive $\tilde{\delta} = 1$, varying R_{max}				
	$R_{max} = 1$	$R_{max} = min(2\tilde{R}; 1)$	$R_{max} = min(1.5\tilde{R}; 1)$	$R_{max} = min(1.25\tilde{R}; 1)$
Share With Adjusted β Same Sign as $\tilde{\beta}$ *Sample: Add Controls, Moves toward Zero*	20%	34%	42%	61%
Share with Adjusted β +/- 2.8 SE of $\tilde{\beta}$ *Sample: All*	13%	33%	36%	46%
Panel B: Randomized Data, Share of Results which Survive $\tilde{\delta} = 1$, varying R_{max}				
	$R_{max} = 1$	$R_{max} = min(2\tilde{R}; 1)$	$R_{max} = min(1.5\tilde{R}; 1)$	$R_{max} = min(1.25\tilde{R}; 1)$
Share With Adjusted β Same Sign as $\tilde{\beta}$ *Sample: Add Controls, Moves toward Zero*	42%	82%	91%	97%
Share with Adjusted β +/- 2.8 SE of $\tilde{\beta}$ *Sample: All*	37%	82%	86%	91%

Notes: This table describes the survival of non-randomized (Panel A) and randomized (Panel B) results under the proportional selection adjustment. Both panels show the share of results which would survive $\tilde{\delta}$ with varying R_{max} values. I consider two definitions of survival: (1) the identified set does not include zero and (2) the outer bound of the set is within 2.8 standard errors of $\tilde{\beta}$. The first of these is considered only for results which move toward zero when controls are added.

Table 5: **Example: Varying R-Squared Movements**

Reference (Table, Columns)	*Baseline Effect (Std.Error)*[R^2]	*Controlled Effect (Std.Error)*[R^2]	*% Change in Coefficient*	*% Change in R-Squared*	*Bias-Adjusted β* $R_{max} = 1.3\tilde{R}, \delta = 1$
Lavy (2009) (Table 4, Math Avg., Column 1 to 2)	5.47 (3.24) [.24]	5.31 (2.73) [.39]	-3.0%	66.4%	5.18
Brunnermeier & Nagel (2008) (Table 5, Column 4 to 5)	-0.108 (.031) [.0879]	-0.103 (.036) [.0881]	-4.2%	0.2%	4.71
Olken & Barron (2009)) (Table 2, Panel B, Column 1 to 2)	-0.735 (.064) [.438]	-0.695 (.071) [.552]	-5.5%	26%	-0.608

Notes: This table shows three examples from the database of economics papers with similar percent change in coefficients and varying R-squared changes. The final column shows the bias-adjusted β with the assumption that $\delta = 1$ and $R_{max} = 1.3\tilde{R}$. The latter is the cutoff I derive from randomized data.

Table 6: **Example: Deviation of Simple and General Estimator**

Result Description	*Baseline Effect (Std.Error)*[R^2]	*Controlled Effect (Std.Error)*[R^2]	*Simple Approx. β* $R_{max} = \tilde{R} + (\tilde{R} - \mathring{R})$	*Simple Approx. β* $R_{max} = 1.3\tilde{R}$	*General Estimator β* $R_{max} = 1.3\tilde{R}$
Trust Relatives	-0.193 (.043) [.106]	-0.178 (.031) [.130]	-0.162	-0.153	0.352
Trust Neighbors	-0.238 (.044) [.115]	-0.202 (.029) [.159]	-0.173	-0.171	-0.044
Trust Local Council	-0.177 (.027) [.175]	-0.128 (.021) [.205]	-0.080	-0.028	0.821
Intragroup Trust	-0.208 (.041) [.121]	-0.187 (.032) [.155]	-0.167	-0.160	0.100
Intergroup Trust	-0.145 (.031) [.093]	-0.115 (.030) [.119]	-0.084	-0.072	0.194

Notes: This table shows the results from Nunn and Wantchekon (2011), Table 4. The first columns show the baseline and controlled effects. The third shows the value of β implied by using the simple approximation and assuming the unobservables add as much to the R-squared as the observables. This is the calculation performed in their paper. The fourth column shows the value of β implied by using the simple approximation and the cutoff derived from the randomized data. The final column shows the correct bias-adjusted β under the cutoff assumptions implied by the randomized data.

Appendix for Online Publication Only

Appendix A: Theoretical Results

Appendix A.1: Single Variable, $\delta = 1$

Proof of Proposition 1

The proof proceeds by calculating each difference in terms of variance and covariance values, and then simplifying.

Claim: $(\mathring{\beta} - \tilde{\beta}) \xrightarrow{p} \frac{\sigma_{1X}}{\sigma_{XX}}(1 - \frac{\sigma_{1X}}{\sigma_{11}}\Pi)$ where $\Pi = \frac{\sigma_{22}\sigma_{1X}}{\sigma_{11}(\sigma_{XX} - \frac{\sigma_{1X}^2}{\sigma_{11}})}$, the asymptotic bias on $\tilde{\beta}$. **Proof:** Observe that $\hat{\lambda}_{w_1^o|X}$ converges in probability to $\frac{Cov(w_1^o,X)}{V(X)} = \frac{\sigma_{1X}}{\gamma_1\sigma_{XX}}$. By a similar logic, $\hat{\lambda}_{W_2|X}$ converges to $\frac{\sigma_{2X}}{\sigma_{XX}}$ and, under proportional selection, to $\frac{\sigma_{1X}\sigma_{22}}{\sigma_{11}\sigma_{XX}}$. $\hat{\lambda}_{W_2|X,w_1^o}$ converges in probability to $\frac{Cov(W_2,X)}{Var(\tilde{X})}$ where $\tilde{X}$ is the residual from a regression of X on $w_{1.}^o$ Note that $Var(\tilde{X})$ converges in probability to $\sigma_{XX} - \frac{\sigma_{1X}^2}{\sigma_{11}}$. Therefore, again invoking proportional selection, $\hat{\lambda}_{W_2|X,w_1^o}$ converges in probability to $\frac{\sigma_{22}\sigma_{1X}}{\sigma_{11}(\sigma_{XX} - \frac{\sigma_{1X}^2}{\sigma_{11}})}$. It will simplify notation to denote this bias Π and to write the equations in terms of Π, σ_{11} and σ_{1X} :

$$\begin{aligned}
\mathring{\beta} &\xrightarrow{p} \beta + \gamma_1\left(\frac{\sigma_{1X}}{\gamma_1\sigma_{XX}}\right) + \Pi\frac{(\sigma_{XX} - \frac{\sigma_{1X}^2}{\sigma_{11}})}{\sigma_{XX}} \\
\tilde{\beta} &\xrightarrow{p} \beta + \Pi
\end{aligned}$$

Subtracting yields:

$$\begin{aligned}
\mathring{\beta} - \tilde{\beta} &\xrightarrow{p} \left(\frac{\sigma_{1X}}{\sigma_{XX}}\right) + \Pi\frac{(\sigma_{XX} - \frac{\sigma_{1X}^2}{\sigma_{11}})}{\sigma_{XX}} - \Pi \\
&\xrightarrow{p} \left(\frac{\sigma_{1X}}{\sigma_{XX}}\right)\left(1 - \frac{\sigma_{1X}}{\sigma_{11}}\Pi\right)
\end{aligned}$$

Claim: $(\tilde{R} - \mathring{R})\hat{\sigma}_{yy} \xrightarrow{p} \sigma_{11} + \Pi^2(\sigma_{XX} - \frac{\sigma_{1X}^2}{\sigma_{11}}) - \frac{1}{\sigma_{XX}}\left(\sigma_{1X} + \Pi(\sigma_{XX} - \frac{\sigma_{1X}^2}{\sigma_{11}})\right)^2$ and $(R_{max} - \tilde{R})\hat{\sigma}_{yy} \xrightarrow{p} \Pi\left(\frac{\sigma_{11}\left(\sigma_{XX} - \frac{\sigma_{1X}^2}{\sigma_{11}}\right)}{\sigma_{1X}} - \Pi\left(\sigma_{XX} - \frac{\sigma_{1X}^2}{\sigma_{11}}\right)\right)$. **Proof:** Observe the following definitions. From the short regression coefficient, $\mathring{R}\hat{\sigma}_{yy} = \sigma_{XX}(\beta + \gamma_1\hat{\lambda}_{w_1^o|X} + \hat{\lambda}_{W_2|X})^2$. By Lemma 1, this converges in probability to $\sigma_{XX}\left(\beta + \left(\frac{\sigma_{1X}}{\sigma_{XX}}\right) + \Pi\frac{(\sigma_{XX} - \frac{\sigma_{1X}^2}{\sigma_{11}})}{\sigma_{XX}}\right)^2$. In the intermediate regression the calculation relies on the coefficient on X $(\beta + \Pi)$ and the coefficient on w_1^o, which is also biased by the exclusion of W_2 through the joint correlation with X and is equal to $\gamma_1 - \frac{\gamma_1\sigma_{1X}}{\sigma_{11}}\Pi$. Note that $\sigma_{11} = \gamma_1^2 Var(w_1^o)$. Thus,
$\tilde{R}\hat{\sigma}_{yy} \xrightarrow{p} \sigma_{XX}(\beta + \Pi)^2 + \sigma_{11}(1 - \frac{\sigma_{1X}}{\sigma_{11}}\Pi)^2 + 2\sigma_{1X}(\beta + \Pi)(1 - \frac{\sigma_{1X}}{\sigma_{11}}\Pi)$. This simplifies to
$\tilde{R}\hat{\sigma}_{yy} \xrightarrow{p} \beta^2\sigma_{XX} + \sigma_{11} + \Pi^2\left(\sigma_{XX} - \frac{\sigma_{1X}^2}{\sigma_{11}}\right) + 2\sigma_{1X}\beta - 2\beta\Pi\left(\sigma_{XX} - \frac{\sigma_{1X}^2}{\sigma_{11}}\right)$. Finally, observe that
$R_{max}\sigma_{yy} = \beta^2\sigma_{XX} + \sigma_{11} + \Pi\frac{\sigma_{11}(\sigma_{XX} - \frac{\sigma_{1X}^2}{\sigma_{11}})}{\sigma_{1X}} + 2\beta\sigma_{1,X} + 2\beta\Pi(\sigma_{XX} - \frac{\sigma_{1X}^2}{\sigma_{11}})$.

Subtracting yields:

$$\begin{aligned}
\left(\tilde{R} - \mathring{R}\right)\hat{\sigma}_{yy} &\xrightarrow{p} \beta^2\sigma_{XX} + \sigma_{11} + \Pi^2\left(\sigma_{XX} - \frac{\sigma_{1X}^2}{\sigma_{11}}\right) + 2\sigma_{1X}\beta - 2\beta\Pi\left(\sigma_{XX} - \frac{\sigma_{1X}^2}{\sigma_{11}}\right) - \sigma_{XX}\left(\beta + \left(\frac{\sigma_{1X}}{\sigma_{XX}}\right) + \Pi\frac{(\sigma_{XX} - \frac{\sigma_{1X}^2}{\sigma_{11}})}{\sigma_{XX}}\right)^2 \\
&\xrightarrow{p} \sigma_{11} + \Pi^2(\sigma_{XX} - \frac{\sigma_{1X}^2}{\sigma_{11}}) - \frac{1}{\sigma_{XX}}\left(\sigma_{1X} + \Pi(\sigma_{XX} - \frac{\sigma_{1X}^2}{\sigma_{11}})\right)^2
\end{aligned}$$

$$
\begin{aligned}
(R_{max}-\tilde{R})\hat{\sigma}_{yy} \quad & \xrightarrow{p} \quad \beta^2\sigma_{XX}+\sigma_{11}+\Pi\frac{\sigma_{11}(\sigma_{XX}-\frac{\sigma_{1X}^2}{\sigma_{11}})}{\delta\sigma_{1X}}+2\beta\sigma_{1,X}+2\beta\Pi(\sigma_{XX}-\frac{\sigma_{1X}^2}{\sigma_{11}}) \\
& \quad -\left(\beta^2\sigma_{XX}+\sigma_{11}+\Pi^2\left(\sigma_{XX}-\frac{\sigma_{1X}^2}{\sigma_{11}}\right)+2\sigma_{1X}\beta-2\beta\Pi\left(\sigma_{XX}-\frac{\sigma_{1X}^2}{\sigma_{11}}\right)\right) \\
& \xrightarrow{p} \quad \Pi\left(\frac{\sigma_{11}\left(\sigma_{XX}-\frac{\sigma_{1X}^2}{\sigma_{11}}\right)}{\sigma_{1X}}-\Pi\left(\sigma_{XX}-\frac{\sigma_{1X}^2}{\sigma_{11}}\right)\right)
\end{aligned}
$$

Claim: Define $\beta^* = \tilde{\beta} - \left[\mathring{\beta} - \tilde{\beta}\right] \frac{R_{max}-\tilde{R}}{\tilde{R}-\mathring{R}}$. Then, $\beta^* \xrightarrow{p} \beta$.
Proof: The claims above define a system of three equations in three unknowns:

$$
\begin{aligned}
(\mathring{\beta}-\tilde{\beta}) \quad & \xrightarrow{p} \quad \left(\frac{\sigma_{1X}}{\sigma_{XX}}\right)\left(1-\frac{\sigma_{1X}}{\sigma_{11}}\Pi\right) \\
\left(\tilde{R}-\mathring{R}\right)\hat{\sigma}_{yy} \quad & \xrightarrow{p} \quad \sigma_{11}+\Pi^2(\sigma_{XX}-\frac{\sigma_{1X}^2}{\sigma_{11}})-\frac{1}{\sigma_{XX}}\left(\sigma_{1X}+\Pi(\sigma_{XX}-\frac{\sigma_{1X}^2}{\sigma_{11}})\right)^2 \\
(R_{max}-\tilde{R})\hat{\sigma}_{yy} \quad & \xrightarrow{p} \quad \Pi\left(\frac{\sigma_{11}\left(\sigma_{XX}-\frac{\sigma_{1X}^2}{\sigma_{11}}\right)}{\sigma_{1X}}-\Pi\left(\sigma_{XX}-\frac{\sigma_{1X}^2}{\sigma_{11}}\right)\right)
\end{aligned}
$$

The unknowns are σ_{11}, σ_{1X} and Π. The system is identified and the solution is $\Pi = \left[\mathring{\beta}-\tilde{\beta}\right]\frac{R_{max}-\tilde{R}}{\tilde{R}-\mathring{R}}$. This gives the result.

Appendix A.2: General Estimator

The approach for the general estimator is virtually identical to the simple estimator case. It will be helpful to observe that $\mathring{\beta}$ and $\mathring{R}$ are defined exactly as above, as is R_{max}.These are not impacted by the way we define the W_1index since the controls are either (1) not included or (b) fully included.
Claim: $(\mathring{\beta}-\tilde{\beta}) \xrightarrow{p} \frac{\sigma_{1X}}{\sigma_{XX}} - \Pi\left(\frac{\sigma_{XX}-\tau_x}{\sigma_{XX}}\right)$ where $\Pi = \frac{\delta\sigma_{1X}\sigma_{22}}{\sigma_{11}\tau_x}$, the asymptotic bias on $\tilde{\beta}$.
Proof: As above, $\mathring{\beta} \xrightarrow{p} \beta + \frac{\sigma_{1X}}{\sigma_{XX}} + \frac{\sigma_{2X}}{\sigma_{XX}}$. Given the proportional selection assumption and the definition of Π above, we therefore have: $\mathring{\beta} \xrightarrow{p} \beta + \frac{\sigma_{1X}}{\sigma_{XX}} + \Pi\frac{\tau_x}{\sigma_{XX}}$. And by definition $\tilde{\beta} \xrightarrow{p} \beta + \Pi$. Differencing yields:

$$
\mathring{\beta}-\tilde{\beta} \quad \xrightarrow{p} \quad \frac{\sigma_{1X}}{\sigma_{XX}} - \Pi\left(\frac{\sigma_{XX}-\tau_x}{\sigma_{XX}}\right)
$$

Claim: $(\tilde{R}-\mathring{R})\hat{\sigma}_{yy} \xrightarrow{p} \sigma_{11}+\Pi^2\tau_x-\frac{1}{\sigma_{XX}}\left(\sigma_{1X}+\Pi\tau_x\right)^2$ and $(R_{max}-\tilde{R})\hat{\sigma}_{yy} \xrightarrow{p} \Pi\left(\frac{\sigma_{11}\tau_x}{\delta\sigma_{1X}}-\Pi\tau_x\right)$.
Proof: By the simple definition of $\mathring{R}$ we have $\mathring{R}\hat{\sigma}_{yy} \xrightarrow{p} (\beta+\frac{\sigma_{1X}}{\sigma_{XX}}+\Pi\frac{\tau_x}{\sigma_{XX}})^2\sigma_{XX}$. Define the variance of element ω_i^o as $\sigma_{ii}^{\omega_o}$ and the covariance of element ω_i^o with X as $\sigma_{ix}^{\omega_o}$. By definition, we have:

$$
\tilde{R}\hat{\sigma}_{yy} \xrightarrow{p} (\beta+\Pi)^2\sigma_{XX} + \sum_{i=1}^{j}\left(\left(\psi_i-\frac{\sigma_{ix}^{\omega_o}}{\sigma_{ii}^{\omega_o}}(\Pi)\right)^2\sigma_{ii}^{\omega_o}\right) + 2(\beta+\Pi)\sum_{i=1}^{j}\left(\sigma_{ix}^{\omega_o}\left(\psi_i-\frac{\sigma_{ix}^{\omega_o}}{\sigma_{ii}^{\omega_o}}(\Pi)\right)\right)
$$

Note the following: $\sigma_{11} = \sum_{i=1}^{j}\psi_i^2\sigma_{ii}^{\omega_o}$, $\sigma_{1X} = \left[\sum_{i=1}^{j}\psi_i\sigma_{ix}^{\omega_o}\right]$ and $\tau_x = \left[\sigma_{XX}-\sum_{i=1}^{j}\frac{\left(\sigma_{ix}^{\omega_o}\right)^2}{\sigma_{ii}^{\omega_o}}\right]$. This therefore simplifies to: $\tilde{R}\hat{\sigma}_{yy} \xrightarrow{p} \beta^2\sigma_{xx}+\sigma_{11}+\Pi^2\tau_x+2\beta\Pi(\tau_x)+2\beta\sigma_{1X}$.

Finally, $R_{max}\hat{\sigma}_{yy} \xrightarrow{p} \beta^2\sigma_{XX}+\sigma_{11}+\sigma_{22}+2\beta\sigma_{1,X}+2\beta\Pi\tau_x$. Subtracting yields both results.
Completion of Proof: The above claims provide a system of three equations in three unknowns (σ_{11},σ_{1X},

Π):

$$
\begin{aligned}
(\mathring{\beta} - \tilde{\beta}) &\xrightarrow{p} \frac{\sigma_{1X}}{\sigma_{XX}} - \Pi \left(\frac{\sigma_{XX} - \tau_x}{\sigma_{XX}} \right) \\
\left(\tilde{R} - \mathring{R} \right) \hat{\sigma}_{yy} &\xrightarrow{p} \sigma_{11} + \Pi^2(\tau_x) - \frac{1}{\sigma_{XX}} \left(\sigma_{1X} + \Pi(\tau_x) \right)^2 \\
(R_{max} - \tilde{R})\hat{\sigma}_{yy} &\xrightarrow{p} \Pi \left(\frac{\sigma_{11}\tau_x}{\sigma_{1X}} - \Pi\tau_x \right)
\end{aligned}
$$

Solving these for Π yields a cubic equation:

$$
\begin{aligned}
0 \quad = \quad & \delta \left((R_{max} - \tilde{R})\sigma_{yy} \right) \left(\mathring{\beta} - \tilde{\beta} \right) \sigma_{XX} \\
+ \quad & \Pi \left(\delta \left((R_{max} - \tilde{R})\sigma_{yy} \right) (\sigma_{XX} - \tau_x) - \left(\left(\tilde{R} - \mathring{R} \right) \sigma_{yy} \right) \tau_x - \sigma_{XX}\tau_x \left(\mathring{\beta} - \tilde{\beta} \right)^2 \right) \\
+ \quad & \Pi^2 \left(\tau_x A \left(\mathring{\beta} - \tilde{\beta} \right) \sigma_{XX} (\delta - 2) \right) \\
+ \quad & \Pi^3 (\delta - 1)(\tau_x \sigma_{XX} - \tau_x^2)
\end{aligned}
$$

The solution will be one of the three roots of this equation (if the equation has three real roots) or the single real root if it has one real root.

Appendix B: Appendix Tables

Table 1: **Summary Statistics: NLSY Wage Data**

	Mean	*Standard Deviation*	Range	*Sample Size*
Log Wages (1996-1998)	2.67	0.63	0-6.21	7496
Years of educ.	12.5	2.24	0-20	7496
Years of exper.	16.3	3.02	8-31	7496
Female	0.49	0.50	0-1	7496
Region of Residence	N/A	N/A	1-4	7496
White	0.64	0.47	0-1	7496
Married Codes	N/A	N/A	0-6	7496
Mother Educ (yrs)	11.0	3.00	1-20	7496
Father Educ (yrs)	11.2	5.27	1-20	7496
Mother Occup (codes)	N/A	N/A	0-984	7496
Father Occup (codes)	N/A	N/A	0-984	7496
Siblings (#)	3.8	2.6	0-22	7496

Notes: This table shows summary statistics for the data used in the NSLY wage analysis in Section 4. Data comes from the NLSY-79 cohort. Means are not reported for region, marital codes or occupation because they are not meaningful. All variables are controlled in the regressions as dummies. Wages are the max of 1996 and 1998 wages.

Table 2: **Summary Statistics: Early Life and Child IQ**

Panel A: IQ Analysis			
	Mean	*Standard Deviation*	*Sample Size*
IQ (PIAT Score, Standardized)	0.025	0.991	6962
Breastfeeding Months	2.40	4.63	6514
LBW + Preterm	0.049	0.217	6174
Mom Drink at all in Pregnancy	0.322	0.467	6537
Age	5.57	1.37	6962
Child Female	0.495	0.500	6962
Mother Black	0.282	0.450	6962
Mother Age	25.3	5.61	6962
Mother Education (years)	12.2	2.7	6962
Mother Income	$41,294	$80,735	6962
Mother Married	0.654	0.476	6962
Panel B: Birth Weight Analysis			
Birth Weight (grams)	3290.4	647.69	7686
Mom Smoke in Pregnancy	0.291	0.454	7686
Drinking Intensity (0-7)	0.638	1.16	7442
Child Female	0.486	0.499	7686
Mother Black	0.273	0.445	7686
Mother Age	24.4	5.49	7686
Mother Education (years)	12.0	2.7	7686
Mother Income	$30,813	$65,374	7686
Mother Married	0.667	0.471	7686

Notes: This table shows summary statistics for the data used in the analysis in Section 4. Drinking intensity is coded from 0 (never) to 7 (every day). Natality detail files are from 2001 and 2002. Data is from the NLSY Children and Young Adults panel.

Appendix C: Alternative Parametrization of R_{max}

This appendix considers how the results in Section 5 would change if I used an alternative parametrization of R_{max}. The primary analysis in the paper uses $R_{max} = \Pi\tilde{R}$ with varying Π. Here, I use $R_{max} = \tilde{R} + \Pi(\tilde{R} - \mathring{R})$. I consider the same questions: the level of robustness of non-randomized results, the Π cutoff implied by the randomized data and what share of non-randomized results would survive the cutoff. For simplicity, I consider here only the primary robustness criteria of whether the identified set excludes zero, and therefore limit to results where inclusion of controls moves the effect size toward zero.

Figure 1a shows the robustness of non-randomized data under this parametrization, and Figure 1b shows the randomized robustness. The multiplier values in both cases are larger here, reflecting the fact that the increase in R-squared from $\mathring{R}$ to $\tilde{R}$ is smaller in value than the level, $\tilde{R}$. The observation that the randomized data is more robust than the non-randomized is even more true. Table 1 replicates the form of Table 4 in the paper. I consider larger values of Π as cutoffs but, again, the conclusion of varying stability and higher stability of randomized results.

The robustness cutoff value implied by the randomized data is 1.9. With this value 47% of non-randomized results would survive, similar to the share in the main analysis.

Table 1: **Robustness of Stability Results, Alternative R_{max}**

Share of Results which Survive $\tilde{\delta} = 1$, varying R_{max}: Robustness is Identified Set Excludes Zero				
	$R_{max} = 1$	$R_{max} = min\{\tilde{R} + 3(\tilde{R} - \mathring{R}), 1\}$	$R_{max} = min\{\tilde{R} + 2(\tilde{R} - \mathring{R}), 1\}$	$R_{max} = min\{\tilde{R} + (\tilde{R} - \mathring{R}), 1\}$
Non-Randomized	20%	37%	44%	66%
Randomized	42%	91%	91%	97%

Notes: This table describes the survival of non-randomized and randomized results under the proportional selection adjustment with varying R_{max} using the alternative R_{max} parametrization. Both rows show the share of results which would survive $\delta = 1$ with varying R_{max} values with survival defined as the identified set does not include zero. The analysis includes only results which move toward zero when controls are added.

Figure 1: **Stability Results Using Additive R_{max} Parametrization**

(a) Rejection of Zero, Non-Randomized, $R_{max} = \tilde{R} + \Pi(\tilde{R} - \mathring{R})$.

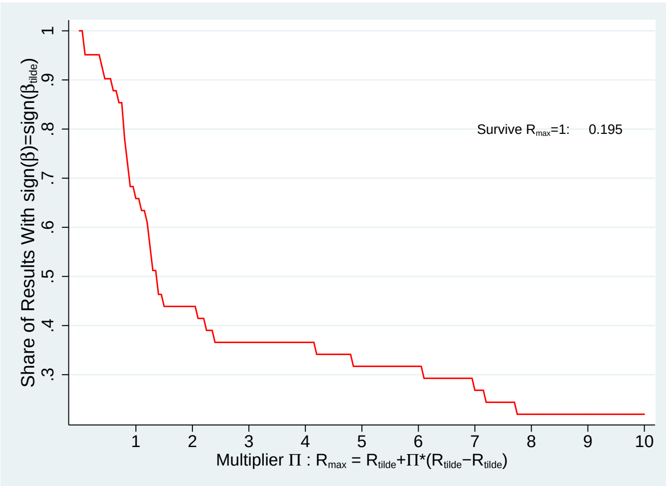

(b) Rejection of Zero, Randomized, $R_{max} = \tilde{R} + \Pi(\tilde{R} - \mathring{R})$.

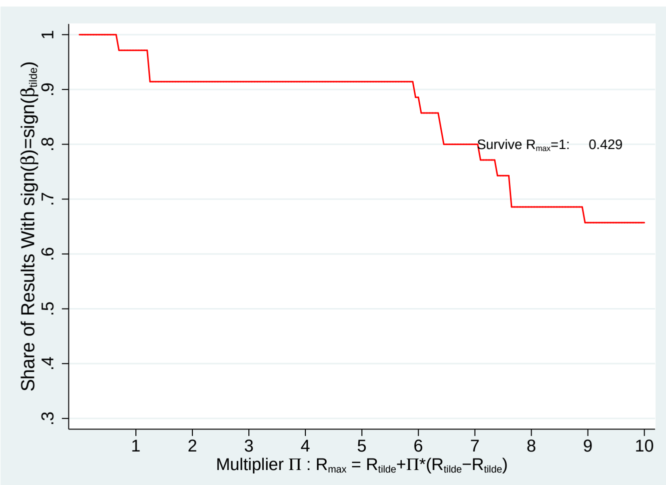

Notes: These graphs show the performance of non-randomized results (Sub-Figure a) and randomized results (Sub-Figure b) with under the proportional selection adjustment. Each figure graphs the share of results which would survive varying parametrizations of R_{max}, in all cases assuming $\tilde{\delta} = 1$. Each Sub-Figure indicates the share of results which would survive $R_{max} = \tilde{R} + \Pi(\tilde{R} - \mathring{R})$ for varying values of Π, with the survival in this case meaning the identified set does not include zero. This figure contains only relationships where the effect is significant with controls and adding the controls moves the coefficient toward zero.

Appendix D: References for Section 6

Abeler, Johannes, Armin Falk, Lorenz Goette, and David Huffman, "Reference Points and Effort Provision," *AMERICAN ECONOMIC REVIEW*, APR 2011, *101* (2), 470–492.

Acemoglu, Daron, Simon Johnson, James A. Robinson, and Pierre Yared, "Income and democracy," *AMERICAN ECONOMIC REVIEW*, JUN 2008, *98* (3), 808–842.

Aghion, Philippe, Robin Burgess, Stephen J. Redding, and Fabrizio Zilibotti, "The Unequal Effects of Liberalization: Evidence from Dismantling the License Raj in India," *AMERICAN ECONOMIC REVIEW*, SEP 2008, *98* (4), 1397–1412.

Aker, Jenny C., Christopher Ksoll, and Travis J. Lybbert, "Can Mobile Phones Improve Learning? Evidence from a Field Experiment in Niger," *AMERICAN ECONOMIC JOURNAL-APPLIED ECONOMICS*, OCT 2012, *4* (4), 94–120.

Algan, Yann and Pierre Cahuc, "Inherited Trust and Growth," *AMERICAN ECONOMIC REVIEW*, DEC 2010, *100* (5), 2060–2092.

Angrist, Joshua and Victor Lavy, "The Effects of High Stakes High School Achievement Awards: Evidence from a Randomized Trial," *AMERICAN ECONOMIC REVIEW*, SEP 2009, *99* (4), 1384–1414.

Ashraf, Nava, "Spousal Control and Intra-Household Decision Making: An Experimental Study in the Philippines," *AMERICAN ECONOMIC REVIEW*, SEP 2009, *99* (4), 1245–1277.

——, **James Berry, and Jesse M. Shapiro**, "Can Higher Prices Stimulate Product Use? Evidence from a Field Experiment in Zambia," *AMERICAN ECONOMIC REVIEW*, DEC 2010, *100* (5), 2383–2413.

Ashraf, Quamrul and Oded Galor, "Dynamics and Stagnation in the Malthusian Epoch," *AMERICAN ECONOMIC REVIEW*, AUG 2011, *101* (5), 2003–2041.

Bandiera, Oriana, Andrea Prat, and Tommaso Valletti, "Active and Passive Waste in Government Spending: Evidence from a Policy Experiment," *AMERICAN ECONOMIC REVIEW*, SEP 2009, *99* (4), 1278–1308.

Beaman, Lori, Raghabendra Chattopadhyay, Esther Duflo, Rohini Pande, and Petia Topalova, "POWERFUL WOMEN: DOES EXPOSURE REDUCE BIAS?," *QUARTERLY JOURNAL OF ECONOMICS*, NOV 2009, *124* (4), 1497–1540.

Beaudry, Paul, David A. Green, and Benjamin Sand, "Does Industrial Composition Matter for Wages? A Test of Search and Bargaining Theory," *ECONOMETRICA*, MAY 2012, *80* (3), 1063–1104.

Becker, Sascha O. and Ludger Woessmann, "WAS WEBER WRONG? A HUMAN CAPITAL THEORY OF PROTESTANT ECONOMIC HISTORY," *QUARTERLY JOURNAL OF ECONOMICS*, MAY 2009, *124* (2), 531–596.

Bernard, Andrew B., Stephen J. Redding, and Peter K. Schott, "Multiple-Product Firms and Product Switching," *AMERICAN ECONOMIC REVIEW*, MAR 2010, *100* (1), 70–97.

Bloom, Nicholas, Raffaella Sadun, and John Van Reenen, "Americans Do IT Better: US Multinationals and the Productivity Miracle," *AMERICAN ECONOMIC REVIEW*, FEB 2012, *102* (1), 167–201.

Bohnet, Iris, Fiona Greig, Benedikt Herrmann, and Richard Zeckhauser, "Betrayal aversion: Evidence from Brazil, China, Oman, Switzerland, Turkey, and the United States," *AMERICAN ECONOMIC REVIEW*, MAR 2008, *98* (1), 294–310.

Boivin, Jean, Marc P. Giannoni, and Ilian Mihov, "Sticky Prices and Monetary Policy: Evidence from Disaggregated US Data," *AMERICAN ECONOMIC REVIEW*, MAR 2009, *99* (1), 350–384.

Broda, Christian, Nuno Limao, and David E. Weinstein, "Optimal Tariffs and Market Power: The Evidence," *AMERICAN ECONOMIC REVIEW*, DEC 2008, *98* (5), 2032–2065.

Brunnermeier, Markus K. and Stefan Nagel, "Do wealth fluctuations generate time-varying risk aversion? Micro-evidence on individuals' asset allocation," *AMERICAN ECONOMIC REVIEW*, JUN 2008, *98* (3), 713–736.

Buera, Francisco J., Alexander Monge-Naranjo, and Giorgio E. Primiceri, "Learning the Wealth of Nations," *ECONOMETRICA*, JAN 2011, *79* (1), 1–45.

Burde, Dana and Leigh L. Linden, "Bringing Education to Afghan Girls: A Randomized Controlled Trial of Village-Based Schools," *AMERICAN ECONOMIC JOURNAL-APPLIED ECONOMICS*, JUL 2013, *5* (3), 27–40.

Bursztyn, Leonardo and Lucas C. Coffman, "The Schooling Decision: Family Preferences, Intergenerational Conflict, and Moral Hazard in the Brazilian Favelas," *JOURNAL OF POLITICAL ECONOMY*, JUN 2012, *120* (3), 359–397.

Bustos, Paula, "Trade Liberalization, Exports, and Technology Upgrading: Evidence on the Impact of MERCOSUR on Argentinian Firms," *AMERICAN ECONOMIC REVIEW*, FEB 2011, *101* (1), 304–340.

Caballero, Ricardo J., Takeo Hoshi, and Anil K. Kashyap, "Zombie Lending and Depressed Restructuring in Japan," *AMERICAN ECONOMIC REVIEW*, DEC 2008, *98* (5), 1943–1977.

Cai, Hongbin, Yuyu Chen, and Hanming Fang, "Observational Learning: Evidence from a Randomized Natural Field Experiment," *AMERICAN ECONOMIC REVIEW*, JUN 2009, *99* (3), 864–882.

Card, David, Alexandre Mas, Enrico Moretti, and Emmanuel Saez, "Inequality at Work: The Effect of Peer Salaries on Job Satisfaction," *AMERICAN ECONOMIC REVIEW*, OCT 2012, *102* (6), 2981–3003.

Carpenter, Jeffrey, Peter Hans Matthews, and John Schirm, "Tournaments and Office Politics: Evidence from a Real Effort Experiment," *AMERICAN ECONOMIC REVIEW*, MAR 2010, *100* (1), 504–517.

Chaney, Eric, "REVOLT ON THE NILE: ECONOMIC SHOCKS, RELIGION, AND POLITICAL POWER," *ECONOMETRICA*, SEP 2013, *81* (5), 2033–2053.

Charles, Kerwin Kofi and Jonathan Guryan, "Prejudice and Wages: An Empirical Assessment of Becker's The Economics of Discrimination," *JOURNAL OF POLITICAL ECONOMY*, OCT 2008, *116* (5), 773–809.

———, Erik Hurst, and Nikolai Roussanov, "CONSPICUOUS CONSUMPTION AND RACE," *QUARTERLY JOURNAL OF ECONOMICS*, MAY 2009, *124* (2), 425–467.

Chen, Roy and Yan Chen, "The Potential of Social Identity for Equilibrium Selection," *AMERICAN ECONOMIC REVIEW*, OCT 2011, *101* (6), 2562–2589.

Chetty, Raj, Adam Looney, and Kory Kroft, "Salience and Taxation: Theory and Evidence," *AMERICAN ECONOMIC REVIEW*, SEP 2009, *99* (4), 1145–1177.

___ **, John N. Friedman, Nathaniel Hilger, Emmanuel Saez, Diane Whitmore Schanzenbach, and Danny Yagan**, "How Does Your Kindergarten Classroom Affect Your Earnings? Evidence from Project Star," *QUARTERLY JOURNAL OF ECONOMICS*, NOV 2011, *126* (4), 1593–1660.

Cohen, Jessica and Pascaline Dupas, "FREE DISTRIBUTION OR COST-SHARING? EVIDENCE FROM A RANDOMIZED MALARIA PREVENTION EXPERIMENT," *QUARTERLY JOURNAL OF ECONOMICS*, FEB 2010, *125* (1), 1–45.

Cohen, Lauren, Andrea Frazzini, and Christopher Malloy, "The Small World of Investing: Board Connections and Mutual Fund Returns," *JOURNAL OF POLITICAL ECONOMY*, OCT 2008, *116* (5), 951–979.

Cole, Shawn, Xavier Gine, Jeremy Tobacman, Petia Topalova, Robert Townsend, and James Vickery, "Barriers to Household Risk Management: Evidence from India," *AMERICAN ECONOMIC JOURNAL-APPLIED ECONOMICS*, JAN 2013, *5* (1), 104–135.

Cortes, Patricia, "The effect of low-skilled immigration on U. S. prices: Evidence from CPI data," *JOURNAL OF POLITICAL ECONOMY*, JUN 2008, *116* (3), 381–422.

Dafny, Leemore, Mark Duggan, and Subramaniam Ramanarayanan, "Paying a Premium on Your Premium? Consolidation in the US Health Insurance Industry," *AMERICAN ECONOMIC REVIEW*, APR 2012, *102* (2), 1161–1185.

Das, Jishnu, Stefan Dercon, James Habyarimana, Pramila Krishnan, Karthik Muralidharan, and Venkatesh Sundararaman, "School Inputs, Household Substitution, and Test Scores," *AMERICAN ECONOMIC JOURNAL-APPLIED ECONOMICS*, APR 2013, *5* (2), 29–57.

Dell, Melissa, "The Persistent Effects of Peru's Mining Mita," *ECONOMETRICA*, NOV 2010, *78* (6), 1863–1903.

Djankov, Simeon, Oliver Hart, Caralee McLiesh, and Andrei Shleifer, "Debt Enforcement around the World," *JOURNAL OF POLITICAL ECONOMY*, DEC 2008, *116* (6), 1105–1150.

Duflo, Esther, Michael Kremer, and Jonathan Robinson, "Nudging Farmers to Use Fertilizer: Theory and Experimental Evidence from Kenya," *AMERICAN ECONOMIC REVIEW*, OCT 2011, *101* (6), 2350–2390.

___ **, Pascaline Dupas, and Michael Kremer**, "Peer Effects, Teacher Incentives, and the Impact of Tracking: Evidence from a Randomized Evaluation in Kenya," *AMERICAN ECONOMIC REVIEW*, AUG 2011, *101* (5), 1739–1774.

___ **, Rema Hanna, and Stephen P. Ryan**, "Incentives Work: Getting Teachers to Come to School," *AMERICAN ECONOMIC REVIEW*, JUN 2012, *102* (4), 1241–1278.

Dupas, Pascaline and Jonathan Robinson, "Why Don't the Poor Save More? Evidence from Health Savings Experiments," *AMERICAN ECONOMIC REVIEW*, JUN 2013, *103* (4), 1138–1171.

Duranton, Gilles and Matthew A. Turner, "The Fundamental Law of Road Congestion: Evidence from US Cities," *AMERICAN ECONOMIC REVIEW*, OCT 2011, *101* (6), 2616–2652.

Estevadeordal, Antoni, Caroline Freund, and Emanuel Ornelas, "DOES REGIONALISM AFFECT TRADE LIBERALIZATION TOWARD NONMEMBERS?," *QUARTERLY JOURNAL OF ECONOMICS*, NOV 2008, *123* (4), 1531–1575.

Fairlie, Robert W. and Jonathan Robinson, "Experimental Evidence on the Effects of Home Computers on Academic Achievement among Schoolchildren," *AMERICAN ECONOMIC JOURNAL-APPLIED ECONOMICS*, JUL 2013, *5* (3), 211–240.

Ferraz, Claudio and Frederico Finan, "Electoral Accountability and Corruption: Evidence from the Audits of Local Governments," *AMERICAN ECONOMIC REVIEW*, JUN 2011, *101* (4), 1274–1311.

Finan, Frederico and Laura Schechter, "Vote-Buying and Reciprocity," *ECONOMETRICA*, MAR 2012, *80* (2), 863–881.

Finkelstein, Amy, "E-ZTAX: TAX SALIENCE AND TAX RATES," *QUARTERLY JOURNAL OF ECONOMICS*, AUG 2009, *124* (3), 969–1010.

Gabaix, Xavier and Augustin Landier, "Why has CEO pay increased so much?," *QUARTERLY JOURNAL OF ECONOMICS*, FEB 2008, *123* (1), 49–100.

Gentzkow, Matthew and Jesse M. Shapiro, "What Drives Media Slant? Evidence From US Daily Newspapers," *ECONOMETRICA*, JAN 2010, *78* (1), 35–71.

Gruber, Jonathan and Daniel M. Hungerman, "The church versus the mall: What happens when religion faces increased secular competition?," *QUARTERLY JOURNAL OF ECONOMICS*, MAY 2008, *123* (2), 831–862.

Guiso, Luigi, Paola Sapienza, and Luigi Zingales, "CULTURAL BIASES IN ECONOMIC EXCHANGE?," *QUARTERLY JOURNAL OF ECONOMICS*, AUG 2009, *124* (3), 1095–1131.

Guryan, Jonathan, Kory Kroft, and Matthew J. Notowidigdo, "Peer Effects in the Workplace: Evidence from Random Groupings in Professional Golf Tournaments," *American Economic Journal: Applied Economics*, October 2009, *1* (4), 34–68.

Heffetz, Ori and Moses Shayo, "How Large Are Non-Budget-Constraint Effects of Prices on Demand?," *AMERICAN ECONOMIC JOURNAL-APPLIED ECONOMICS*, OCT 2009, *1* (4), 170–199.

Ifcher, John and Homa Zarghamee, "Happiness and Time Preference: The Effect of Positive Affect in a Random-Assignment Experiment," *AMERICAN ECONOMIC REVIEW*, DEC 2011, *101* (7), 3109–3129.

Jayachandran, Seema and Adriana Lleras-Muney, "LIFE EXPECTANCY AND HUMAN CAPITAL INVESTMENTS: EVIDENCE FROM MATERNAL MORTALITY DECLINES," *QUARTERLY JOURNAL OF ECONOMICS*, FEB 2009, *124* (1), 349–397.

Jensen, Robert, "Do Labor Market Opportunities Affect Young Women's Work and Family Decisions? Experimental Evidence from India," *QUARTERLY JOURNAL OF ECONOMICS*, MAY 2012, *127* (2), 753–792.

Jensen, Robert T. and Nolan H. Miller, "Giffen Behavior and Subsistence Consumption," *AMERICAN ECONOMIC REVIEW*, SEP 2008, *98* (4), 1553–1577.

Karlan, Dean S. and Jonathan Zinman, "Credit elasticities in less-developed economies: Implications for microfinance," *AMERICAN ECONOMIC REVIEW*, JUN 2008, *98* (3), 1040–1068.

Khwaja, Asim Ijaz and Atif Mian, "Tracing the Impact of Bank Liquidity Shocks: Evidence from an Emerging Market," *AMERICAN ECONOMIC REVIEW*, SEP 2008, *98* (4), 1413–1442.

Kirwan, Barrett E., "The Incidence of US Agricultural Subsidies on Farmland Rental Rates," *JOURNAL OF POLITICAL ECONOMY*, FEB 2009, *117* (1), 138–164.

Kling, Jeffrey R., Sendhil Mullainathan, Eldar Shafir, Lee C. Vermeulen, and Marian V. Wrobel, "Comparison Friction: Experimental Evidence from Medicare Drug Plans," *QUARTERLY JOURNAL OF ECONOMICS*, FEB 2012, *127* (1), 199–235.

Kremer, Michael, Jessica Leino, Edward Miguel, and Alix Peterson Zwane, "SPRING CLEANING: RURAL WATER IMPACTS, VALUATION, AND PROPERTY RIGHTS INSTITUTIONS," *QUARTERLY JOURNAL OF ECONOMICS*, FEB 2011, *126* (1), 145–205.

Kroft, Kory, Fabian Lange, and Matthew J. Notowidigdo, "Duration Dependence and Labor Market Conditions: Evidence from a Field Experiment," *QUARTERLY JOURNAL OF ECONOMICS*, AUG 2013, *128* (3), 1123–1167.

Lalive, Rafael and Josef Zweimueller, "HOW DOES PARENTAL LEAVE AFFECT FERTILITY AND RETURN TO WORK? EVIDENCE FROM TWO NATURAL EXPERIMENTS," *QUARTERLY JOURNAL OF ECONOMICS*, AUG 2009, *124* (3), 1363–1402.

Lavy, Victor, "Performance Pay and Teachers' Effort, Productivity, and Grading Ethics," *AMERICAN ECONOMIC REVIEW*, DEC 2009, *99* (5), 1979–2011.

Libecap, Gary D. and Dean Lueck, "The Demarcation of Land and the Role of Coordinating Property Institutions," *JOURNAL OF POLITICAL ECONOMY*, JUN 2011, *119* (3), 426–467.

Mas, Alexandre and Enrico Moretti, "Peers at Work," *AMERICAN ECONOMIC REVIEW*, MAR 2009, *99* (1), 112–145.

Mian, Atif and Amir Sufi, "THE CONSEQUENCES OF MORTGAGE CREDIT EXPANSION: EVIDENCE FROM THE US MORTGAGE DEFAULT CRISIS," *QUARTERLY JOURNAL OF ECONOMICS*, NOV 2009, *124* (4), 1449–1496.

Michalopoulos, Stelios and Elias Papaioannou, "PRE-COLONIAL ETHNIC INSTITUTIONS AND CONTEMPORARY AFRICAN DEVELOPMENT," *ECONOMETRICA*, JAN 2013, *81* (1), 113–152.

Muralidharan, Karthik and Venkatesh Sundararaman, "Teacher Performance Pay: Experimental Evidence from India," *JOURNAL OF POLITICAL ECONOMY*, FEB 2011, *119* (1), 39–77.

Nunn, Nathan, "The long-term effects of Africa's slave trades," *QUARTERLY JOURNAL OF ECONOMICS*, FEB 2008, *123* (1), 139–176.

——— **and Leonard Wantchekon**, "The Slave Trade and the Origins of Mistrust in Africa," *AMERICAN ECONOMIC REVIEW*, DEC 2011, *101* (7), 3221–3252.

Olken, Benjamin A. and Patrick Barron, "The Simple Economics of Extortion: Evidence from Trucking in Aceh," *JOURNAL OF POLITICAL ECONOMY*, JUN 2009, *117* (3), 417–452.

Pope, Devin G. and Maurice E. Schweitzer, "Is Tiger Woods Loss Averse? Persistent Bias in the Face of Experience, Competition, and High Stakes," *AMERICAN ECONOMIC REVIEW*, FEB 2011, *101* (1), 129–157.

Price, Joseph and Justin Wolfers, "RACIAL DISCRIMINATION AMONG NBA REFEREES," *QUARTERLY JOURNAL OF ECONOMICS*, NOV 2010, *125* (4), 1859–1887.

Rockoff, Jonah E., Douglas O. Staiger, Thomas J. Kane, and Eric S. Taylor, "Information and Employee Evaluation: Evidence from a Randomized Intervention in Public Schools," *AMERICAN ECONOMIC REVIEW*, DEC 2012, *102* (7), 3184–3213.

Schularick, Moritz and Alan M. Taylor, "Credit Booms Gone Bust: Monetary Policy, Leverage Cycles, and Financial Crises, 1870-2008," *AMERICAN ECONOMIC REVIEW*, APR 2012, *102* (2), 1029–1061.

Shayo, Moses and Asaf Zussman, "Judicial Ingroup Bias in the Shadow of Terrorism," *QUARTERLY JOURNAL OF ECONOMICS*, AUG 2011, *126* (3), 1447–1484.

Snyder, Jr. James M. and David Stromberg, "Press Coverage and Political Accountability," *JOURNAL OF POLITICAL ECONOMY*, APR 2010, *118* (2), 355–408.

Song, Zheng, Kjetil Storesletten, and Fabrizio Zilibotti, "Growing Like China," *AMERICAN ECONOMIC REVIEW*, FEB 2011, *101* (1), 196–233.

Spolaore, Enrico and Romain Wacziarg, "THE DIFFUSION OF DEVELOPMENT," *QUARTERLY JOURNAL OF ECONOMICS*, MAY 2009, *124* (2), 469–529.

Sullivan, Daniel and Till von Wachter, “JOB DISPLACEMENT AND MORTALITY: AN ANALYSIS USING ADMINISTRATIVE DATA,” *QUARTERLY JOURNAL OF ECONOMICS*, AUG 2009, *124* (3), 1265–1306.

Voors, Maarten J., Eleonora E. M. Nillesen, Philip Verwimp, Erwin H. Bulte, Robert Lensink, and Daan P. Van Soest, “Violent Conflict and Behavior: A Field Experiment in Burundi,” *AMERICAN ECONOMIC REVIEW*, APR 2012, *102* (2), 941–964.

Wei, Shang-Jin and Xiaobo Zhang, “The Competitive Saving Motive: Evidence from Rising Sex Ratios and Savings Rates in China,” *JOURNAL OF POLITICAL ECONOMY*, JUN 2011, *119* (3), 511–564.