

Credible Inference for the Returns to Schooling

Jonathan L. Gu

May 22, 2020

Abstract

This article estimates the heterogeneous returns to education by place-of-birth for males born between 1930 and 1940. I contrast the traditional instrumental variables estimates that rely on features on the asymptotic distribution with finite-sample Bayesian analysis which relies on a prior distribution for the uncertain parameters. The traditional IV literature has recently been supplemented with tests for "weak instruments" that are intended to determine if the asymptotic results can be relied on, while the Bayesian advice is: it depends on the prior distribution. Traditional Bayesian analysis tries to duck the problem of choosing a prior distribution by using a uniform prior over the whole line to represent "ignorance" but when the target is the inverse of a parameter (the denominator of an IV estimator) a uniform prior for the parameter is distinctly not uniform for its inverse. When using priors determined from cross-validation, we are 95% sure that the returns to education are positive for only four out of nine regions, whereas the standard IV-2SLS asymptotic distribution would yield "significant" results for all nine geographic regions.

Keywords: Machine Learning, Causality, Bayesian, Weak Instrumental Variables, Heterogeneous Returns to Schooling, Angrist-Krueger data.

Contents

1	Introduction	3
2	Data	4
3	Comparison of Instrumental Variables Methods	10
3.1	Setup of Simple Model	10
3.2	Gibbs sampling from Posterior of Simple Model	13
3.3	Comparison of Methods	16
4	Model for Heterogeneous Returns to Schooling	17
4.1	Instrumental Variables Assumptions	18
4.2	Augmented First-Stage Reduced Form	18
4.3	Error Structure	20
4.4	Priors	21
5	Estimation	23
5.1	Conditional Posterior for the Covariance Structure	24
5.2	Conditional Posterior for the Coefficients	24
6	Results	27
7	Conclusion	29

1 Introduction

There is a long history of using instrumental variables to determine the causal effect of schooling on education. Angrist and Krueger (1991) finds that each additional year of schooling can boost earnings by about 8%. The discussion of methods for determining our confidence in these point estimates has spurred a lot of literature. Current empirical applications of instrumental variables face must first pass a binary test of whether the instruments are "strong" enough for us to assume that the likelihood function is quadratic, and therefore the estimate is distributed normal with variance as derived from the delta method. This article is the first to estimate and make statements of confidence about the heterogeneous returns to education by region of birth using the same 1980 Census data used by Angrist and Krueger. We do not suffer from "weak instruments" by avoiding using an asymptotic approximation (Leamer, 2010) of the instrumental variables two-stage least squares (IV-2SLS) estimate.

Instead, we sample from the posterior distribution of our two-stage estimate and form point estimates and measures of confidence directly from the samples. When using weak priors, we obtain the same point estimates as the standard IV-2SLS methods, but we have much larger 95% credible intervals. When using the "best" priors as determined from cross-validation, we find that we are 95% sure that the returns to education are positive for only four out of nine regions, whereas the standard IV-2SLS approach would yield "significant" results for all nine geographic regions.

Bound et al. (1995) uses simulated (fake) instruments to obtain similar results of "significance" to the results obtained by (Angrist and Krueger, 1991). In the case of IV-2SLS, standard definitions of significance is defined by whether zero lies in a 95% confidence interval formed by applying the delta method to the two-stage least squares estimate. The delta method is an appeal to the large sample asymptotics of the IV-2SLS estimate. However, we know that the just-identified IV-2SLS estimate does not have finite moments (Nelson and Startz, 1990; Phillips, 2009)¹, and therefore the appeal to asymptotics in the case of IV-2SLS analysis is diminished.

Chamberlain and Imbens (1996) explains that sampling from the posterior instead of using an asymptotic approximation leads to tight posterior intervals when using the data, and wide posterior intervals when using randomly simulated instruments. Lopes and Polson (2014) and Hoogerheide et al. (2007) discuss a variety of priors to further analyze the causal returns to education in the same setting. This article is the first to add the analysis of heterogeneous treatment effects to the tradition of applying Bayesian credible

¹The number of finite moments depends on the degree of over-identification.

inference methods towards empirical data.

Due to the increased availability of big data and large scale AB tests in the tech industry, readily available code packages (EconML, 2019) inspired by (Newey and Powell, 2003; Hartford et al., 2017) have surfaced that advertise the ability to determine heterogeneous treatment effects in a nonparametric manner. For the initial impact of the AB test on an initial outcome, the method for discovering causal heterogeneous effects is quite straightforward. However, once we ask for the causal effect of the initial outcome on a secondary outcome, we begin to confront the issue of weak instruments. Staiger and Stock (1994); Stock and Yogo (2002) discuss conditions under which the instrument is "strong enough" to assume the IV-2SLS estimate is normal.

Their approach requires the source of heterogeneity to be non-stochastic, or independent from the omitted factors that determine the treatment. In the case of heterogeneous returns to schooling by region-of-birth, this means that we require the region-of-birth to be independent from the omitted variables that determine the years of schooling. Instead, we implement a strategy that avoids this pitfall at the cost of including estimating additional first-stage equations.


Section 2 summarizes the data, Section 3 compares using the posterior distribution with standard instrumental variables research design, Section 4 sets up a model of heterogeneous treatment effects and shows why we have many first-stage equations, Section 5 discusses the method of sampling from the posterior, Section 6 summarizes our results, and Section 7 concludes.

2 Data

Following (Angrist and Keueger, 1991), we use the 5 percent Public Use sample (the A Sample) of the US population as of April 1, 1980. Following Chamberlain and Imbens (1996), we filtered down to native-born African American/Black or Caucasian/white males with birthdays in the first or fourth quarter of a year between 1930 and 1939 (inclusive). The key insight for our identification strategy is that men born in the fourth quarter tend to have about one more year of schooling. Figure 1 shows the interquartile ranges of the years of education against quarter-of-birth. For census participants born in the fourth quarter, we see that the first quartile of the level of schooling is equal to the median. This suggests there is a binding lower bound to the number of years of education for census participants born in the fourth quarter.

I limit the sample to men with positive wage and salary earnings, and positive weeks worked in 1979, leaving 202,859 participants in the sample. We compute weekly earn-

Figure 1: Quarter of Birth Against Years of Education



This figure shows the interquartile ranges of the years of education against quarter-of-birth. For census participants born in the fourth quarter, we see that the first quartile of the level of schooling is equal to the median. This suggests there is a binding lower bound to the number of years of education for census participants born in the fourth quarter.

ings by dividing annual earnings by weeks worked. [Table 1](#) shows that the average man was 45 years old and earned \$430 dollars per week in 1979. Eighty four percent of these men were married, and nine percent of these men were African American/Black. The level of schooling is determined by the years of completed schooling, where twelve corresponds to completing high school. The original paper from 1991 conditioned on some concomitant variables such as place of current residence, and marital status, however conditioning on concomitant variables can corrupt the causal interpretation of the data, so I remove these controls. For this article, we weight each census respondent equally, and we leave extensions to weighted calculations for future research.

Table 1: Descriptive Statistics

	Mean	SD	10th Pct	90th Pct
Weekly Wages	427.4837	253.9709	179.4828	692.4039
Years of Education	12.68296	3.301034	8	17
Age	44.90577	2.948244	41	49
Married	.8406331	.3660187	0	1
African America	.0946618	.292748	0	0

N = 202,859 individuals in dataset.

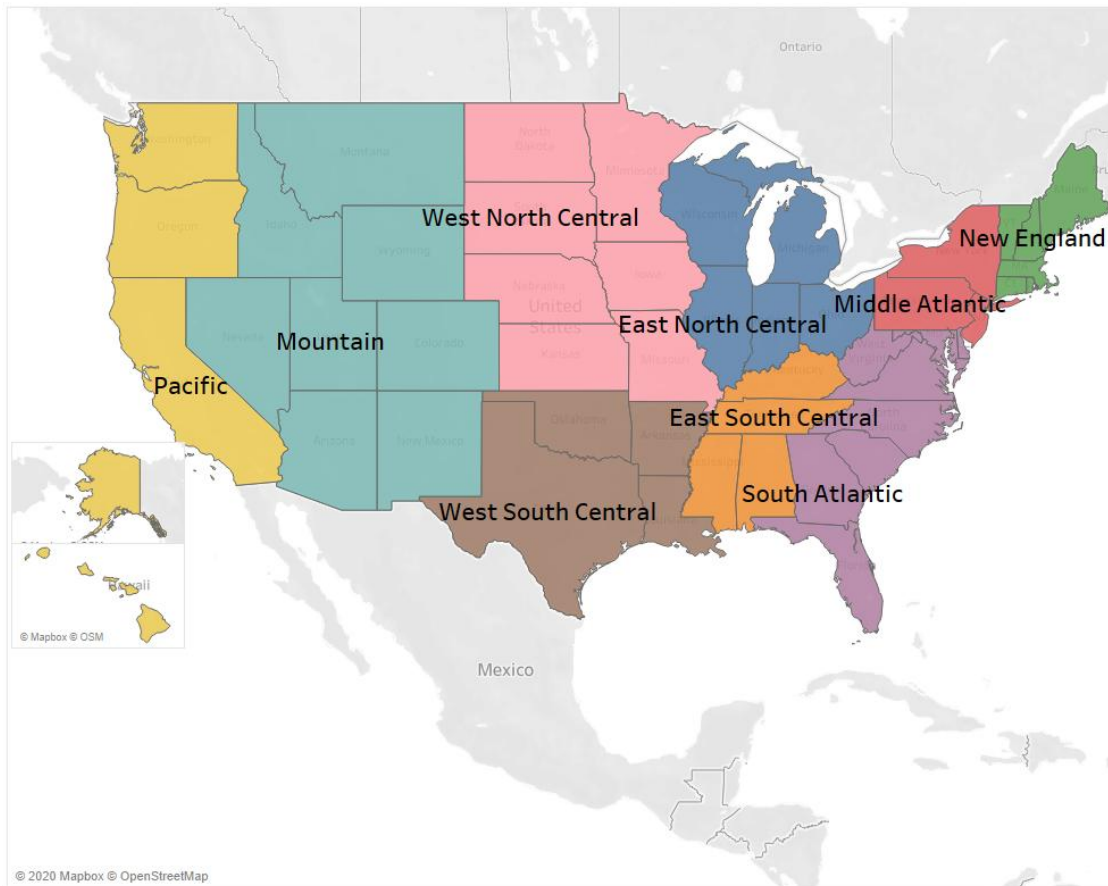
Source: 1980 Census (5% Public Use Sample A). We filtered down to native born African American/Black or Caucasian/white males with birthdays in the first or fourth quarter of a year between 1930 and 1939 (inclusive) with positive wage and salary earnings, and positive weeks worked in 1979.

Since some of the states are much smaller than the others, we follow the original paper in using census defined geographic regions as the source of potentially heterogeneous treatment effects. Since the region of birth is determined chronologically before any educational and employment decisions we can worry less about the issue of concomitant variables. [Figure 2](#) displays the census regions. [Table 2](#) lists out the states that make up each region.

[Figure 3](#) shows the relationship between the average log weekly earnings against the average years of education for each census region. We see a general positive trend, but more importantly, we see that each region has differing average levels of weekly earnings and education. This article checks if the returns to education differs across these geographic regions.

The quarter-of-birth is our instrument. We follow the original paper in creating more instruments by interacting the quarter of birth with the census regions and the years of birth. Adding all these instruments would make us worry about adding "weak" instruments, which would negatively affect inference in the IV-2SLS estimation approach because weak instruments make the IV-2SLS more biased and further away from being nor-

Figure 2: Census Regions

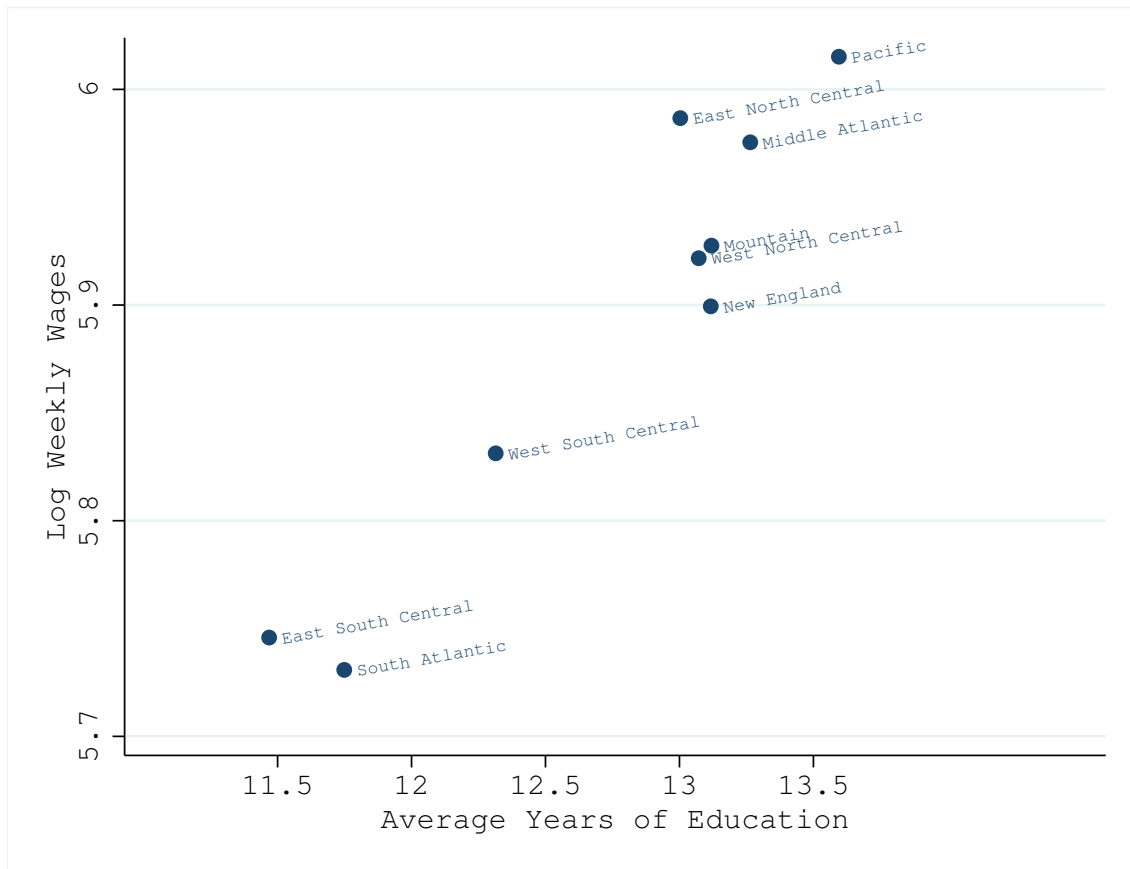


This figure shows the census defined geographic regions. [Table 2](#) lists out the states that make up each region.

Table 2: Geographic Regions

<u>Census Region</u>	<u>Birth State</u>	<u>Census Region</u>	<u>Birth State</u>
East North Central	Illinois Indiana Michigan Ohio Wisconsin	West North Central	Iowa Kansas Minnesota Missouri Nebraska North Dakota South Dakota
East South Central	Alabama Kentucky Mississippi Tennessee	West South Central	Arkansas Louisiana Oklahoma Texas
Mountain	Arizona Colorado Idaho Montana Nevada New Mexico Utah Wyoming	Middle Atlantic	New Jersey New York Pennsylvania
South Atlantic	Delaware District Of Columbia Florida Georgia Maryland North Carolina South Carolina Virginia West Virginia	New England	Connecticut Maine Massachusetts New Hampshire Rhode Island Vermont
		Pacific	Alaska California Hawaii Oregon Washington

Figure 3: Income Against Education by Census Region



This figure shows the relationship between the average log weekly earnings against the average years of education for each census region. We see a general positive trend, but more importantly, we see that each region has differing average levels of weekly earnings and education. This article checks if the returns to education differs across these geographic regions.

mally distributed. Due to this issue, we don't assume that the final estimate is distributed normally. Instead we sample from the posterior of the our instrumental variables estimate.

3 Comparison of Instrumental Variables Methods

This section compares the standard Instrumental Variables approach to a one one that examines the actual posterior dsitribution of a Bayesian estimate. To easily compare with previous literature, we focus on the nationwide returns to schooling first (instead of the heterogeneous returns to schooling). The current rule of thumb for for when one "can" use the instrumental variables is that the first stage F-statistic should be large – greater than 20. This section confirms that a confidence interval obtained from using the actual posterior distribution of the IV-2SLS approach agrees with one obtained from using the asymptotic distribution. We also plug in a variety of priors using estimates of the re- turns to education obtained from other studies and other datasets, by using these other estimates as priors, we essentially pool the data from the two datasets. Of course, this is assuming that the two datasets are independent, and that the returns to education are homogeneous.

3.1 Setup of Simple Model

We start of with a simple model of the returns to education. We assume that the returns to education are homogeneous for now:

$$Y_i = S_i\beta_S + X_i\beta_{SX} + \beta_{Sc} + \varepsilon_i \quad (1)$$

$$S_i = Z_i\pi_S + X_i\pi_{SX} + \pi_{Sc} + \eta_i \quad (2)$$

Where Y_i is weekly earnings in 1979, S_i is the level of schooling, and our instrument, Z_i , is an indicator of whether the student was born in the fourth quarter. The controls X_i include the year of birth, birthplace, and race.

We can also remove the projection of the data onto X . Let \perp refer to the residuals after projecting onto the controls and the constant:

$$Y_i^\perp = S_i^\perp\beta_S + \varepsilon_i^\perp \quad (3)$$

$$S_i^\perp = Z_i^\perp\pi_S + \eta_i^\perp \quad (4)$$

The simple instrumental variables assumptions are as follows:

Assumption 1. *The instrument is relevant, excluded, and unconfounded:*

$$\text{Relevance: } \pi_S \neq 0 \quad (5)$$

$$\text{Exclusion: } (Z_i^\perp \perp Y_i^\perp) | S_i \quad (6)$$

$$\text{Unconfoundedness: } Z_i^\perp \perp (\varepsilon_i, \eta_i) \quad (7)$$

Finally, we assume that the errors are distributed normally:

Assumption 2. *The errors are multivariate normal with mean zero and variance Σ :*

$$(\varepsilon_i, \eta_i) \sim N((0, 0), \Sigma)$$

The IV2SLS estimate is the value of β_S that maximizes the likelihood of observing our data. Recall that by the Frish-Waugh-Lovell theorem, we can simply use the data after removing the projection onto X to obtain our point estimates. Let the subscript N denote a column of stacked data for each individual:

$$\tilde{\pi}_{2SLS} = (Z_N^\perp Z_N^\perp)^{-1} (Z_N^\perp S_N^\perp) \quad (8)$$

$$\hat{S}_N^\perp = Z_N^\perp \tilde{\pi}_{2SLS} \quad (9)$$

$$\tilde{\beta}_{2SLS} = (\hat{S}_N^\perp \hat{S}_N^\perp)^{-1} (\hat{S}_N^\perp Y_N^\perp) \quad (10)$$

According to the standard research design for IV-2SLS estimates, [Staiger and Stock \(1994\)](#); [Stock and Yogo \(2002\)](#) propose using the first stage F-test for whether π is zero:

$$F_Z \equiv (N - 1) \frac{SSE_{null} - SSE_Z}{SSE_Z} \quad (11)$$

$$SSE_Z = \hat{\eta}_N^\perp \hat{\eta}_N^\perp \quad (12)$$

$$SSE_{null} = S_N^\perp S_N^\perp \quad (13)$$

$$\hat{\eta}_N^\perp = S_N^\perp - \hat{S}_N^\perp \quad (14)$$

The current standard in instrumental variables analysis is to use the asymptotic distribution of the estimate if F_Z is greater 10 ([Staiger and Stock, 1994](#)) or 20 ([Stock and Yogo, 2002](#)). The asymptotic distribution of the IV-2SLS estimate is normal, centered at the true value, and with variance derived from taking the inverse of the Fisher information matrix (or equivalently, from using the delta method):

$$\widehat{\text{Var}}(\tilde{\beta}_{2\text{SLS}}) = (\widehat{\text{RH}}_{2\text{N}}' \widehat{\text{RH}}_{2\text{N}})^{-1} \hat{\eta}'_{\text{N}} \hat{\eta}_{\text{N}} / (\text{N} - 1) \quad (15)$$

$$\widehat{\text{RH}}2 = [\text{Z}\tilde{\pi}_{2\text{SLS}}, \text{X}, 1] \quad (16)$$

$$\hat{\eta}_{\text{N}} = \text{Y}_{\text{N}} - \widehat{\text{RH}}_{2\text{N}} \tilde{\beta} \quad (17)$$

Where $\widehat{\text{RH}}2$ denotes the right-hand side of the second stage with the predicted levels of schooling plugged in for the actual level of schooling, and $\tilde{\beta} \equiv (\widehat{\text{RH}}2' \widehat{\text{RH}}2)^{-1} \widehat{\text{RH}}2' \text{Y}_{\text{N}}$ is the full set of coefficients in the regression that includes the predicted levels of schooling.

The discussion on weak instruments, and when it is appropriate to use the asymptotic distribution is far from over, see (Young, 2019; Andrews et al., 2019; Bekker, 1994; Hahn and Hausman, 2002) for history of instrumental variables estimates and the variety of methods to determining confidence. The main takeaway is that Equation 15 can yield variances that can be many orders of magnitude too small. In the just identified case (one treatment and one instrument), the IV-2SLS estimate doesn't even have finite variance, therefore the variance as derived from the delta method is infinitely too small, and it is inappropriate to use bootstrap to estimate an infinite variance.

Instead of assuming the asymptotic distribution, we can look at the posterior distribution of our estimate (Lopes and Polson, 2014; Hoogerheide et al., 2007; Chamberlain and Imbens, 1996). To do so, we put priors on the covariance of the errors Σ and the coefficients β, π :

Prior 1.

$$\Sigma \sim \text{IW}(\nu_0 = 0.01, \Sigma_0 = (0.01) * \mathbf{I}_{(\text{K}_G+1)}) \quad (18)$$

$$(\beta, \pi) \sim \text{N}((\mu_\beta, \mu_\pi), \begin{pmatrix} \lambda_\beta^{-1} & 0 \\ 0 & \lambda_\pi^{-1} \end{pmatrix}) \quad (19)$$

Where μ_β and μ_π are vectors that represent our prior point estimates for the values of the coefficients β and π , and λ_β^{-1} and λ_π^{-1} are diagonal matrices that represent the strength of our priors. Throughout this exercise, we keep all prior variances at 1 million, except for the prior on the key coefficient that represents how much schooling affects earnings.

The reduced form set of simultaneous equations is:

$$Y_i = (Z_i\pi_S + X_i\pi_{SX} + \pi_{Sc} + \eta_i)\beta_S + X_i\beta_{SX} + \beta_{Sc} + \varepsilon_i \quad (20)$$

$$= Z_i\pi_S\beta_S + X_i(\pi_{SX}\beta_S + \beta_{SX}) + \pi_{Sc}\beta_S + \beta_{Sc} + \underbrace{(\eta_i\beta_S + \varepsilon_i)}_{\eta_{y,i}} \quad (21)$$

$$= Z_i\pi_S\beta_S + X_i(\pi_{SX}\beta_S + \beta_{SX}) + \pi_{Sc}\beta_S + \beta_{Sc} + \eta_{y,i} \quad (22)$$

$$S_i = Z_i\pi_S + X_i\pi_{SX} + \pi_{Sc} + \eta_i \quad (23)$$

$$\eta_{y,i} \equiv (\eta_i\beta_S + \varepsilon_i) \quad (24)$$

Therefore the reduced form errors $(\eta_{y,i}, \eta_i)$ have the reduced form covariance matrix:

$$\Omega = \begin{pmatrix} 1 & \beta_S \\ 0 & 1 \end{pmatrix} \Sigma \begin{pmatrix} 1 & 0 \\ \beta_S & 1 \end{pmatrix} \quad (25)$$

$$= \begin{pmatrix} \omega_{yy} & \omega_{ys} \\ \omega_{sy} & \omega_{ss} \end{pmatrix} \quad (26)$$

3.2 Gibbs sampling from Posterior of Simple Model

In order to sample from the posterior distribution of $\tilde{\beta}_{2SLS}$, we express the conditional distribution of the three key parameters β_S, π_S, Σ on each other:

$$(\Sigma|\beta_S, \pi_S, Y_N, S_N, Z_N, X_N) \sim IW(n_0 + N, \Sigma_0 + N\hat{\Sigma}_N) \quad (27)$$

$$(\beta_S|\Sigma, \pi_S, Y_N, S_N, Z_N, X_N) \sim N(\mathbf{b}_S, B_S) \quad (28)$$

$$(\pi_S|\Sigma, \beta_S, Y_N, S_N, Z_N, X_N) \sim N(\mathbf{p}_S, P_S) \quad (29)$$

$$\hat{\Sigma}_N = [\hat{\varepsilon}_N, \hat{\eta}_N]'[\hat{\varepsilon}_N, \hat{\eta}_N]$$

$$B_S^{-1} = \lambda_\beta + \widetilde{RH}_{2N}'\widetilde{RH}_{2N}$$

$$(B_S)^{-1}\mathbf{b}_S = \lambda_\beta\mu_\beta + \widetilde{RH}_{2N}'\widetilde{Y}_N$$

$$P_S^{-1} = \lambda_\pi + \widetilde{RH}_{1N}'\widetilde{RH}_{1N}$$

$$(P_S)^{-1}\mathbf{p} = \lambda_\pi\mu_\pi + \widetilde{RH}_{1N}'\widetilde{S}_N$$

$$\widetilde{Y}_N = (Y_N - (\hat{\eta}_N)\omega_{ss}^{-1}\omega_{sy})\omega_{yy|s}^{-0.5}$$

$$\widetilde{S}_N = (S_N - (\hat{\eta}_{y,N})\omega_{yy}^{-1}\omega_{ys})\omega_{ss|y}^{-0.5}$$

$$\widetilde{RH}_{2N} = \widehat{RH}_{2N}\omega_{yy|s}^{-0.5}$$

$$\widetilde{RH}_{1N} = RH_{1N}\omega_{ss|y}^{-0.5}$$

$$\omega_{yy|s} = \omega_{yy} - \omega_{ys}\omega_{ss}^{-1}\omega_{sy}$$

$$\omega_{ss|y} = \omega_{ss} - \omega_{sy}\omega_{yy}^{-1}\omega_{ys}$$

$$\hat{\eta}_N = S_N - RH_{1N}\pi_S$$

$$\hat{\varepsilon}_N = Y_N - RH_{2N}\beta_S$$

$$\hat{\eta}_{y,N} = Y_N - \widehat{RH}_{2N}\beta_S$$

$$RH_{1N} = [Z_N, X_N, 1]$$

$$RH_{2N} = [S_N, X_N, 1]$$

$$\widehat{RH}_{2N} = [RH_{1N}\pi, X_N, 1]$$

Please see [Section 5](#) for a detailed explanation and intuition behind the calculations here. These conditional posteriors are derived by multiplying the likelihood by the prior, and then noticing that the resulting functional form takes the shape of Inverse Wishart or Normal distributions respectively. In order to draw from the posterior distribution of our instrumental variables estimate, we perform Gibbs sampling by iteratively drawing from the distributions expressed in Equations (27, 28, 29).

Table 3: Comparison of Methods of Inference for Instrumental Variables

Standard Research Design					
		First Stage F-Statistic	Point Estimate	95% Confidence Interval Based on Asymptotic Distribution	
				Lower Bound	Upper Bound
[1]	OLS		0.0633	0.0624	0.0641
[2]	IV-2SLS	110.5070	0.0685	0.0356	0.1013

Bayesian Posterior Calculation with Five Different Priors						
		Prior Distribution is Normal with ...		Posterior Distribution		
		Mean	Standard Deviation	Mean	95% Credible Interval	
					Lower Bound	Upper Bound
[3]	-	0	1000	0.0684	0.0354	0.1018
[4]	-	0	0.05	0.0612	0.0292	0.0931
[5]	-	0.15	0.05	0.0765	0.0457	0.1080
[6]	<i>Bettinger et al. (2019)</i> Data: All CA students eligible for Cal Grant (1998 to 2000). IV: Discontinuity in GPA eligibility for Cal Grant.	0.054	0.019	0.0618	0.0367	0.0864
[7]	<i>Card (1995b)</i> Data: NLSY (1966 Cohort) IV: Indicator for living near a 4-year college.	0.097	0.048	0.0715	0.0407	0.1028

Notes:

[1-2]: These point estimates, first stage F-statistic, and 95% confidence intervals are what you would find in a standard instrumental variables research design. The first-stage F-statistic is large enough for us to assume the IV-2SLS estimate is distributed according to its asymptotic distribution.

[3-7]: These point estimates and 95% credible intervals are obtained from sampling from the posterior distribution of the IV-2SLS estimate. The five estimates only vary the prior on the coefficient that determines the impact of schooling on earnings. All other coefficients use a normal prior with mean zero and variance one million.

[3]: Since a prior with such a high variance is relatively uninformative, the point estimate and the 95% credible interval agree with [2].

[4-5]: If we had a relatively strong prior belief that there were no returns to education, then this prior combined with the evidence would lead us to reject our prior point estimate at a 95% credible level. Likewise, if we had a relatively strong prior belief that each year of schooling increased earnings by 15%, then the posterior would lead us to reject our prior point estimate at a 95% credible level.

[6-7]: We use the point estimates and standard deviations reported by other papers. Assuming the returns to education are homogeneous, then the posterior point estimates and credible intervals represent pooling evidence from the two datasets. Since the previous papers' point estimates all lie within the 95% posterior credible interval, the 1980 census data of men born between 1930 and 1940 does not lead us to reject the estimates found in previous literature at a 95% credible level.

3.3 Comparison of Methods

Table 3 summarizes the variety of methods of determining confidence in the instrumental variables estimates. The first two rows show point estimates, first stage F-statistic, and 95% confidence intervals as one would find in a standard instrumental variables research design. The first-stage F-statistic is large enough for us to assume the IV-2SLS estimate is distributed according to its asymptotic distribution. The OLS reports a much tighter confidence interval, but without an identification strategy the estimate does not have a causal interpretation.

The final four rows shows point estimates and 95% credible intervals obtained from sampling from the posterior distribution of the IV-2SLS estimate. The five estimates only vary the prior on the coefficient that determines the impact of schooling on earnings. All other coefficients use a normal prior with mean zero and variance one million.

The third row reports nearly identical point estimates and 95% credible intervals to the standard IV-2SLS approach because we use a weak prior for the returns to education. This supports the claim from (Staiger and Stock, 1994; Stock and Yogo, 2002) that a large enough F-statistic allows us to use the asymptotic distribution of the IV-2SLS estimate. When we examine the heterogeneous returns to education in the next section – with multiple treatments – the asymptotic distribution and the posterior distribution of the IV-2SLS estimates will not agree anymore.

The fourth and fifth row show priors that would have their point estimates rejected at a 95% level when combined with the evidence. If we had a relatively strong prior belief that there were no returns to education, then this prior combined with the evidence would lead us to reject our prior point estimate at a 95% credible level. Likewise, if we had a relatively strong prior belief that each year of schooling increased earnings by 15%, then the posterior would lead us to reject our prior point estimate at a 95% credible level.

Instead of using a weak prior for the returns to schooling, the last three rows use the point estimates and standard deviations reported by other papers. Assuming the returns to education are homogeneous, then the posterior point estimates and credible intervals represents pooling evidence from the two datasets. Since the previous papers' point estimates all lie within the 95% posterior credible interval, we can say that the 1980 census data of men born between 1930 and 1940 does not lead us to reject the estimates found in previous literature at a 95% credible level.

The last two rows use the point estimates and standard deviations reported by other papers. Assuming the returns to education are homogeneous, then the posterior point estimates and credible intervals represent pooling evidence from the two datasets. Since

the previous papers' point estimates all lie within the 95% posterior credible interval, the 1980 census data of men born between 1930 and 1940 does not lead us to reject the estimates found in previous literature at a 95% credible level.

4 Model for Heterogeneous Returns to Schooling

We would like to estimate the heterogeneous impact of schooling on log weekly earnings. The heterogeneity is with respect to the region of birth². Throughout this paper, data for each individual, i , is in row form, and N in the subscript denotes vertically stacked data for all the participants.

$$Y_i = S_i\beta_S + S_iG_i\beta_{SG} + G_i\beta_G + X_i\beta_X + \beta_c + \varepsilon_i \quad (30)$$

Y_i is log weekly earnings, S_i is years of schooling, G_i is a vector of indicators for each region of birth, excluding one to prevent over-saturation. We have a $K_G = 9$ geographic regions. The controls X_i include race and indicators for each age between 1930 and 1940. We have ten controls after leaving out one age indicator and one race indicator to prevent over-saturating the model. Let $K_X = 10$ represent the number of additional controls.

The instrument for schooling is an indicator of whether the respondent was born in the fourth quarter. We have the following reduced form relationships between schooling and the instrument:

$$S_i = Z_i\pi_{S1} + G_i\pi_{G1} + X_i\pi_{X1} + \pi_{c,1} + \eta_{i1} \quad (31)$$

$$S_iG_{k,i} = Z_i\pi_{Sk} + G_i\pi_{Gk} + X_i\pi_{Xk} + \pi_{c,k} + \eta_{ik} \text{ for } k \in \{2, \dots, K_G\} \quad (32)$$

Where Z_i is a $1 \times K_Z$ vector of instruments, which includes any interactions between the quarter-of-birth and any other variables. In our analysis, $K_Z = 19$, which represents interacting the quarter-of-birth with region, age, and race indicators. This is the same approach taken in the original paper. [Bound et al. \(1995\)](#) delves into the issue of many and weak instruments to conclude that it is not appropriate to assume the resulting two-stage least squares estimate is normally distributed due to the issue of weak instruments. [Chamberlain and Imbens \(1996\)](#) shows how we can overcome this issue by examining the posterior distribution instead. This article is an extension of the posterior oriented

²For region definitions, see [Figure 3](#). We leave the extension to random treatment effects for future research.

analysis that examines heterogeneous treatment effects.

4.1 Instrumental Variables Assumptions

We have the usual identifying assumptions:

Assumption 3. *The instrument is relevant, excluded, and unconfounded:*

$$\text{Relevance: } (\pi_S' \pi_S) \text{ is invertible} \quad (33)$$

$$\text{Exclusion: } (Z_i \perp\!\!\!\perp Y_i) | S_i, X_i \quad (34)$$

$$\text{Unconfoundedness: } Z_i \perp\!\!\!\perp (\varepsilon_i, \eta_i) | X_i \quad (35)$$

Where π_S is the matrix of coefficients that predict the treatments from the instruments. Some sufficient conditions for $(\pi_S' \pi_S)$ to be invertible are that $KZ > KG$ (we have more instruments than treatments), and that $\pi_{S,k}$ are all linearly independent (π_S is full column rank). This second sufficient condition means that the set of instruments should impact each treatment differently, otherwise we wouldn't be able to identify the impact two two different treatments from this set of instruments.

Note that Assumptions (34,35) are conditional on the control variables X_i , which can make the assumptions fore tenable: controlling for the birthyear and the birthplace, the location is independent from omitted variables that determine schooling and earnings.

4.2 Augmented First-Stage Reduced Form

We have augmented the initial reduced form relationship (Equation 31) with expressions for each interaction between quarter-of-birth and geographic region (Equation 32). This is a departure from widely available "nonparametric heterogeneous instrumental variables" implementations in the Python package EconML (2019). To see why, we only plug one

equation (Equation 31) into Equation 30:

$$Y_i = (Z_i\pi_{S1} + G_i\pi_{G1} + X_i\pi_{X1} + \pi_{c,1} + \eta_{i1})\beta_{S1} + (Z_i\pi_{S1} + G_i\pi_{G1} + X_i\pi_{X1} + \pi_{c,1} + \eta_{i1})G_i\beta_{SG} + G_i\beta_G + X_i\beta_X + \beta_c + \varepsilon_i \quad (36)$$

$$= Z_i(\pi_{S1}\beta_{S1} + \pi_{S1}G_i\beta_{SG}) + \quad (37)$$

$$G_i(\pi_{G1}\beta_{S1} + \pi_{G1}G_i\beta_{SG} + \beta_G) + X_i(\pi_{X1}\beta_{S1} + \pi_{X1}G_i\beta_{SG} + \beta_X) + \beta_c + \pi_{c,1}(\beta_{S1}) + \pi_{c,1}G_i\beta_{SG} + \eta_{i1}\beta_{S1} + \underbrace{\eta_{i1}G_i}_{\text{interaction}}(\beta_{SG}) + \varepsilon_i \quad (38)$$

Due to the interaction between the first stage reduced form errors and the geographic region indicators in Equation 38, the two-stage least squares estimate will still be biased unless we assume the source of heterogeneity, G_i , is a fixed, non-stochastic vector or if we assume G_i is independent from the reduced form errors. This is like saying that any omitted factors that determine the the years of schooling for each person are independent from the place of birth!

An alternative is to instrument for the heterogeneous treatments also, which is why we include first stage reduced form representations of all the interactions between the amount of education with geography. To see this, let's first compactify the notation by representing the treatments by letting S_i represent all the interacted treatment effects: $S_i \equiv [S_i \quad S_i G_i]$. Then we can stack the coefficients on the treatments β_S and the first stage coefficients on the instruments π_S :

$$S_i = Z_i\pi_S + G_i\pi_G + X_i\pi_X + \pi_c + \eta_i \quad (39)$$

$$Y_i = S_i\beta_S + G_i\beta_G + X_i\beta_X + \beta_c + \varepsilon_i \quad (40)$$

$$= (Z_i\pi_S + G_i\pi_G + X_i\pi_X + \pi_c + \eta_i)\beta_S + G_i\beta_G + X_i\beta_X + \beta_c + \varepsilon_i \\ = Z_i\pi_S\beta_S + G_i(\pi_G\beta_S + \beta_G) + X_i(\pi_X\beta_S + \beta_X) + (\pi_c\beta_S + \beta_c) + \underbrace{\eta_i\beta_S}_{\text{no interaction}} + \varepsilon_i \quad (41)$$

Equation 39 is the *first stage*, Equation 40 is the *second stage*, and Equation 41 is the *second stage reduced form*. The standard two-stage least squares estimation proceeds by predicting values for schooling, then regressing log earnings on these predicted schooling levels. Note that the first stage estimation consists of estimating a reduced form specifi-

cation for each heterogenous treatment effect.

Another way to arrive at the same estimate is to regress *both* the outcomes *and* the treatments on the instruments, then "divide" the two coefficients. In matrix notation "dividing" the two coefficients involves inverting the matrix $(\pi'\pi)$. We have $\beta = E((\pi'\pi)^{-1}\pi'(\pi_{YonZ}))$. Where π_{YonZ} is the regression of Y on the right hand side from the *first-stage*. This is why the relevance condition in [Assumption 3](#) is about the invertibility of a matrix.

In the second stage reduced form ([Equation 41](#)), the first stage errors (η_i) do not interact with the geographic indicators. This is the key reason why we choose to estimate additional first stage parameters – we don't need to assume G_i is non-stochastic, or that G_i and η_i are independent. This also simplifies both the computational and analytical calculation of the conditional posterior distributions because we don't have to deal with heterogeneous reduced form variance-covariance matrices.

By estimating more parameters in the first stage, we drastically raise the threshold of "strength of the first stage" under which we would believe that our final two-stage least squares estimate is distributed normally. [Staiger and Stock \(1994\)](#); [Stock and Yogo \(2002\)](#) display how to check the F-statistic of the first stage for conditions when we can assume the two-stage least squares estimate is asymptotically normal. [Sanderson and Windmeijer \(2016\)](#) carries out the same large-sample-but-worsening-instruments analysis for the case with multiple endogenous variables. As noted by many authors ([Nelson and Startz, 1990](#); [Bekker, 1994](#); [Phillips, 2009](#)) the actual exact distribution of the two stage least squares estimate is not normal. In the just-identified case the exact distribution of the two-stage least squares estimate doesn't even have finite moments, which would negate the use of bootstrap procedures to estimate the mean.

Instead, we work with the posterior distribution of the two-stage least squares estimate and obtain estimates of confidence from the Bayesian posterior distribution. We follow homogeneous treatment effect examples ([Lopes and Polson, 2014](#); [Wiesenfarth et al., 2014](#)) in choosing conjugate priors that are less informative about estimate. Conjugate priors allow us to easily sample from the posterior distribution, save computational time, and write out simple expressions for the conditional posterior distributions.

4.3 Error Structure

Assumption for the Errors/Omitted Variables:

We assume that the errors are normal, and specify priors for the covariances and the coefficients.

Assumption 4. *The errors are multivariate normal with mean zero and variance Σ :*

$$\begin{pmatrix} \varepsilon_i \\ \eta_i \end{pmatrix}_{\substack{1 \times 1 \\ 1 \times (K_G)}} \sim N(\vec{0}, \Sigma)$$

Assuming that the errors are normal is not as strong as assuming the final two-stage least squares estimate is distributed normally. As an example, recall that the most basic two-stage least squares estimate in the single treatment, single instrument case reduces to the ratio of two estimated coefficients. If the errors in the first and second stage are normal, then the two-stage least squares estimate is the ratio of two correlated normal random variables. [Fieller \(1932\)](#) calculates the exact probability density function of the ratio of two correlated normal random variables. In the just-identified case the two-stage least squares estimate doesn't have finite moments ([Nelson and Startz, 1990](#)).

4.4 Priors

Prior for the covariance:

We assume the variance Σ is distributed Inverse Wishart (the conjugate prior), with relatively uninformative choices for its parameters:

Prior 2. *Our prior for Σ is Inverse Wishart:*

$$\Sigma \sim IW(n_0 = 0.01, \Sigma_0 = (0.01) * \mathbf{I}_{(K_G+1)})$$

Where \mathbf{I} represents the identity matrix, and recall that K_G is the number of geographic regions. Choosing small values for n_0 and Σ_0 are uninformative because if we were to observe N additional data: $e_N = (\varepsilon_i, \eta_i)_{i=\{1, \dots, N\}}$, then the posterior distribution for the variance would be $(\Sigma | e_N) \sim IW(n_0 + N, \Sigma_0 + N \hat{\Sigma}_N)$, where $\hat{\Sigma}_N$ is an estimate of the covariance from the data. Therefore, the small initial values for the shape parameters allow the data have a stronger impact on the posterior distribution of Σ .

We can also represent the variance-covariance matrix for the errors in reduced form:

$$\begin{aligned}
\Omega &= \mathbf{B}\Sigma\mathbf{B}' \\
&= \begin{pmatrix} 1 & \beta_S' \\ \vec{0} & \mathbf{I}_{K_G} \end{pmatrix} \Sigma \begin{pmatrix} 1 & \vec{0}' \\ \beta_S & \mathbf{I}_{K_G} \end{pmatrix} \\
&= \begin{pmatrix} 1 & \beta_S' \\ \vec{0} & \mathbf{I}_{K_G} \end{pmatrix} \begin{pmatrix} \sigma_{\varepsilon\varepsilon} & \Sigma_{\varepsilon\eta} \\ \Sigma_{\eta\varepsilon} & \Sigma_{\eta\eta} \end{pmatrix} \begin{pmatrix} 1 & \vec{0}' \\ \beta_S & \mathbf{I}_{K_G} \end{pmatrix} \\
&= \begin{pmatrix} \sigma_{\varepsilon\varepsilon} + 2\Sigma_{\varepsilon\eta}\beta_S + \beta_S'\Sigma_{\eta\eta}\beta_S & \Sigma_{\varepsilon\eta} + \beta_S'\Sigma_{\eta\eta} \\ \Sigma_{\eta\varepsilon} + \Sigma_{\eta\eta}\beta_S & \Sigma_{\eta\eta} \end{pmatrix} \quad (42) \\
&= \begin{pmatrix} \omega_{\eta_0\eta_0} & \Omega_{\eta_0\eta} \\ \Omega_{\eta\eta_0} & \Omega_{\eta\eta} \end{pmatrix} \quad (43)
\end{aligned}$$

Where I have expanded the variance-covariance matrices of the errors into convenient block form, and η_0 is the error from the reduced form regression of Y_i on the right hand side from the first stage. We have a choice of forming our prior for the covariance matrix in terms of Σ or Ω . [Chao and Phillips \(1998\)](#); [Dreze \(1976\)](#) discuss the 1-to-1 correspondence between priors on Σ and priors on Ω . Our current prior puts an identity matrix as the initial shape of the Inverse-Wishart distribution, which corresponds to an initial shape of $\mathbf{B}\mathbf{B}'$ on the prior of the reduced form errors.

Prior for the coefficients:

I use a normal prior for the coefficients:

Prior 3. *Our prior for the coefficients is Multivariate Normal:*

$$(\beta, \pi) \sim N(\vec{0}, \text{Diag}([\lambda_B, \lambda_P])^{-1}) \quad (44)$$

Where (β, π) are the set of coefficients in the structural model, and $\text{Diag}(v)$ corresponds to a matrix with the vector v on the diagonal and zero elsewhere. The main advantage of using a normal prior is because it is the conjugate prior to the specification that the errors are normal.

The normal prior with a diagonal covariance matrix corresponds to using ridge regression with the regularization penalties scaled by λ_B and λ_P . Recall that ridge regression minimizes the squared residuals plus a weighted sum of squares of the coefficients. If we were to simply use ridge regression of Y on X with constant penalties λ , then the analytical solution is: $\beta_{X, \text{ridgereg}} = (X'X + \mathbf{I}\lambda)^{-1}X'Y$, which is the standard analytical OLS solution

with the λ added to the "denominator"³. Taking the regularization penalties towards zero is the same as increasing the variance of the normal priors to infinity, and this will lead us to the standard OLS analytical solution. Throughout this article, I refer to a "weak" prior as one that has variance set to 1 million, setting the ridge regression regularization penalties to $\frac{1}{(1 \text{ million})}$

Even without a Bayesian interpretation, some amount of regularization is still justified as the dimension of the treatment increases. If we were to attempt to estimate nonparametric returns to schooling, we would need to use regularization even if we did not want to have a Bayesian interpretation of the results because nonparametric instrumental variables are plagued by an "ill-posed inverse" problem. The estimation becomes the issue of regressing earnings on an infinite set of functions of our treatment (schooling). Unfortunately, our instrument doesn't impact all possible functions of the treatment. [Newey and Powell \(2003\)](#) shows that a nonparametric instrumental variables strategy can be consistent when one conducts ridge regression for the outcome on the predicted levels of schooling (the second stage), and reformulates ridge regression as a constrained maximization problem, for which the constraint on the magnitude of the coefficients is assumed to not bind.

5 Estimation

In order to obtain estimates of our confidence in the parameters, we sample from the conditional posterior distributions of each parameter (this is Gibbs sampling). Since all of our priors are conjugate distributions, the conditional posteriors take the same shape. First we have the posterior distributions of the variance-covariance matrices:

³This is why ridge regression is also called a shrinkage estimator.

5.1 Conditional Posterior for the Covariance Structure

$$(\Sigma|\beta, \pi, RH_{1N}, RH_{2N}, Y_N) \sim IW(\nu_0 + N, \Sigma_0 + N\hat{\Sigma}_N), \quad (45)$$

$$(\Omega|\Sigma, \beta, \pi, RH_{1N}, RH_{2N}, Y_N) = \mathbf{B}\Sigma\mathbf{B}' \quad (46)$$

$$RH_{2N} \equiv \begin{bmatrix} \mathbf{S}_N & \mathbf{G}_N & \mathbf{X}_N & \mathbf{1} \end{bmatrix} \quad (47)$$

$N \times (2K_G + K_X) \quad N \times K_G \quad N \times (K_G - 1) \quad N \times K_X$

$$RH_{1N} \equiv \begin{bmatrix} \mathbf{Z}_N & \mathbf{G}_N & \mathbf{X}_N & \mathbf{1} \end{bmatrix} \quad (48)$$

$N \times (K_Z + K_G + K_X) \quad N \times K_Z \quad N \times (K_G - 1) \quad N \times K_X$

$$\hat{\Sigma}_N = [\hat{\varepsilon}_N, \hat{\eta}_N]'[\hat{\varepsilon}_N, \hat{\eta}_N] \quad (49)$$

$$\hat{\varepsilon}_N = Y_N - (RH)_{2N}\beta \quad (50)$$

$$\hat{\eta}_N = Y_N - (RH)_{1N}\pi \quad (51)$$

Where Y_N is the column of log weekly earnings and RH_{1N} and RH_{2N} are stacked rows of right hand side variables in the first and second stages respectively. We calculate Ω at each step of the Gibbs sampling process because this will make subsequent draws from the conditional posteriors of β and π much easier. As noted above, $\hat{\Sigma}_N$ is the estimated covariance matrix where $[\hat{\varepsilon}_N, \hat{\eta}_N]$ is an $N \times 2$ matrix of the two vectors of predicted errors horizontally stacked together.

5.2 Conditional Posterior for the Coefficients

Since we have chosen conjugate priors, the posterior distribution for β is also normal:

$$(\beta|\pi, \Sigma, \Omega, RH_{1N}, \widehat{RH}_{2N}, Y_N) \sim N(b_1, B_1) \quad (52)$$

$$\widehat{RH}_{2N} \equiv [RH_{1N}\pi, G_N, X_N, 1]$$

Now we need to specify what the parameters b_1, B_1 are. Note that the posterior for β only conditions on predicted schooling because we want use the variation induced in schooling by differences in the quarter-of-birth to identify β . This is the same reasoning as in any instrumental variables estimation process.

The posterior parameters for β are:

$$B_1^{-1} = \text{Diag}(\lambda_B) + \widetilde{\widehat{RH}}_{2N}' \widetilde{\widehat{RH}}_{2N} \quad (53)$$

$$B_1^{-1} b_1 = \widetilde{\widehat{RH}}_{2N}' \widetilde{Y}_N \quad (54)$$

$$\widetilde{\widehat{RH}}_{2N} = \widehat{RH}_{2N} \omega_{\eta_0 \eta_0 | \eta}^{-0.5} \quad (55)$$

$$\widetilde{Y}_N = (Y_N - \underbrace{(\mathbf{S}_N - \widehat{RH}_{1N} \pi)}_{\hat{\eta}_N} \Omega_{\eta\eta}^{-1} \Omega_{\eta\eta_0}) \omega_{\eta_0 \eta_0 | \eta}^{-0.5} \quad (56)$$

$$\omega_{\eta_0 \eta_0 | \eta} = (\omega_{\eta_0 \eta_0} - \Omega_{\eta_0 \eta} \Omega_{\eta\eta}^{-1} \Omega_{\eta\eta_0}) \quad (57)$$

The term $(\mathbf{S}_N - \widehat{RH}_{1N} \pi)$ our estimate of the reduced form errors.

Sampling for π looks quite similar to sampling from β , however we should break π into the coefficients for their individual first stage equations. Recall from Equations (31,32) that we have one first stage equation for each interaction between the schooling and the geographic indicator because we want to avoid having our first stage reduced form errors interacting with the source of heterogeneity. Then we have K_G first stage equations, one for each geographic region, and π is a matrix of coefficients with dimension $(KZ + KG + KX) \times KG$. Let π_k represent a single column of π , then $\pi = [\pi_1, \pi_2, \dots, \pi_{K_G}]$. Let $\pi_{(-k)}$ represent all the columns of π except for π_k . We can sample from the conditional posterior of π_k conditioning on $\pi_{(-k)}, \beta, \Omega, Y, \widehat{RH}_{1N}, \widehat{RH}_{2N}$:

$$(\pi_k | \pi_{(-k)}, \beta, \Sigma, \Omega, \widehat{RH}_{1N}, \widehat{RH}_{2N}, Y_N) \sim N(b_1, B_1) \quad (58)$$

Since $\pi_{(KZ+KG+KX) \times KG}$ is a matrix of parameters, it will be easier to sample from the posterior distributions of each individual treatment. Let $\pi_k_{(KZ+KG+KX) \times 1}$ represent the k th column of π , and let π_{-k} represent the other columns. For example, π_1 represents the first stage coefficients that predict the first element of \mathbf{S}_i from \widehat{RH}_1 .

Now it is easier to express the conditional posterior distribution of the one dimensional vector π_k .

$$(\pi_k | \pi_{-k}, \beta, \Sigma, \Omega, \widehat{RH}_{1N}, \widehat{RH}_{2N}, Y_N) \sim N(p_k, P_k) \quad (59)$$

We can express the posterior distribution of π are

$$P_k^{-1} = \text{Diag}(\lambda_{p_k}) + \widetilde{R}H'_{1kN} \widetilde{R}H_{1kN} \quad (60)$$

$$P_k^{-1} p_k = \widetilde{R}H'_{1kN} \widetilde{S}_{kN} \quad (61)$$

$$\widetilde{R}H_{1kN} = RH_{1N} \omega_{\eta_k \eta_k | \eta_{(-k)}}^{-0.5} \quad (62)$$

$$\widetilde{S}_{kN} = (S_k - \underbrace{([Y_N, \mathbf{S}_{(-k)}] - [\widehat{R}H_{2N} \beta, RH_{1N} \pi_{(-k)}])}_{[\hat{\eta}_0, \hat{\eta}_{-k}]} \Omega_{\eta_{(-k)} \eta_{(-k)}}^{-1} \Omega_{\eta_{(-k)} \eta_k}) \omega_{\eta_k \eta_k | \eta_{(-k)}}^{-0.5} \quad (63)$$

$$\omega_{\eta_k \eta_k | \eta_{-k}} = \omega_{\eta_k \eta_k} - \Omega_{\eta_k \eta_{(-k)}} \Omega_{\eta_{(-k)} \eta_{(-k)}}^{-1} \Omega_{\eta_{(-k)} \eta_k} \quad (64)$$

Where $\mathbf{S}_{(-k)}$ are the treatments, leaving out the k th treatment, and $\omega_{\eta_k \eta_k}$ refers to the variance of the error in the k th first stage estimating equation. Since all the errors come from the same set of simultaneous equations, it makes sense that the posteriors take the exact same form. Once again, the term $([Y_N, \mathbf{S}_{(-k)}] - [\widehat{R}H_{2N} \beta, RH_{1N} \pi_{(-k)}])$ represents our estimates of the error from the reduced form second stage, and the errors from the first stages excluding the k th error.

The Gibbs sampling procedure using the Julia programming language. I drew 10,000 samples from the Markov Chain Monte Carlo procedure, using the 2SLS estimates as the initial values and using 2,000 initial burn-in simulations to allow the draws to stabilize. Although MCMC samples produces autocorrelated, we do not perform any thinning (only taking every n th observation), as measures of percentiles, medians, means, and variances are more precise without thinning (MacEachern and Berliner, 1994; Link and Eaton, 2012).

When we use relatively uninformative priors, we should obtain the same point estimates as the standard 2SLS approach. However, there is a literature on allowing the data to inform us about about the priors (Zellner et al., 2014; Zellner, 1978). Following recent methods Hartford et al. (2017), we use a k -fold cross validation determine the optimal variances in the normal prior (or the optimal regularization parameters.) K -fold cross validation consists of splitting the data into a training and validation set, then training the model on the training set and evaluates the mean squared error of the predicted outcomes on the validation set. I use Julia to discover the regularization penalties that minimize the out-of-sample mean-squared error. Recall that ridge regression regularization penalties correspond to variances for a normal prior.

6 Results

Table 4 summarizes the results of our exercise. We discuss each column from left to right. The first two columns are point estimates and 95% confidence intervals obtained by using standard techniques. All results are "significant". The OLS results indicate that there is a positive empirical relationship between education and earnings, although this relationship cannot be interpreted as causal. The positive point estimates for the IV estimates indicate that there might be a positive causal relationship between education and earnings. A standard interpretation of the 95% confidence interval as derived from the asymptotic approximation of the distribution of the IV-2SLS estimate would lead us to believe that the returns to education are "significantly" positive for all Census Regions.

The third column is the mean and 95% credible interval are gathered by Gibbs sampling 10,000 times from the posterior of the instrumental variables estimate. We use a "weak" prior that assumes the coefficients have variance 1 million (the ridge regression penalty is $1/(1 \text{ million})$). The point estimates are the mean of the Gibbs samples, while the 95% credible intervals have endpoints determined by the 2.5 and 97.5 percentiles of the samples. Note that the point estimates for the case with weak priors are quite similar to the "standard" IV point estimates, but the 95% credible intervals are so much wider that only one out of nine intervals doesn't include zero. The wider intervals lead us to conclude that it is inappropriate to assume that the "standard" 2SLS estimate is normally distributed with the variance as derived from the delta method.

The final set of estimates is also has a mean and 95% credible interval obtained by Gibbs sampling 10,000 times from the posterior of the instrumental variables estimate. We use K-Fold cross validation to determine that a prior variance of 4.45 for the second stage and upwards of 1 million for the first stage performed the best at out-of-sample prediction (minimizing mean squared error using a two-stage approach).

The rightmost values display the percent of the posterior that lies above zero. This is a measure of our confidence that the returns to education are positive. We conclude that our identification strategy leads us to be 95% confident that the returns to education are positive for the following four census regions: East North Central, East South Central, South Atlantic, West South Central. **Figure 4a** shows a map of the regions and corresponding proportion of the posterior that lies above zero, representing the posterior chance that the returns to education are positive. Our identification strategy leads us to be 95% confident that the returns to education are positive for the following four census regions: East North Central, East South Central, South Atlantic, West South Central.

Table 4: Results

<u>Census Region</u>	<u>Standard Methods</u>		<u>Bayesian Instrumental Variables</u>	
	<u>OLS Coefficient</u> ^[1]	<u>IV Coefficient</u> ^[1]	<u>Posterior (Weak Prior)</u> ^[2]	<u>Posterior (Best Prior)</u> ^{[3][4]}
East North Central	0.059 (0.057,0.0611)	0.0259 (0.0239,0.028)	0.0262 (-0.1072,0.1583)	0.0715 (-0.0114,0.1562) 95.6%
East South Central	0.0606 (0.0581,0.0631)	0.0454 (0.0429,0.0479)	0.0452 (-0.0009,0.0898)	0.0453 (0.0013,0.0889) 97.8%
Middle Atlantic	0.0715 (0.0696,0.0734)	0.0364 (0.0344,0.0384)	0.0348 (-0.1656,0.2326)	0.0403 (-0.1188,0.2001) 69.5%
Mountain	0.0596 (0.0553,0.0639)	0.0245 (0.02,0.0289)	0.0235 (-0.1377,0.1862)	0.0288 (-0.1124,0.1704) 65.8%
New England	0.0687 (0.0654,0.072)	0.0532 (0.0498,0.0565)	0.0499 (-0.0754,0.173)	0.0504 (-0.0626,0.1645) 80.7%
Pacific	0.0542 (0.0504,0.0581)	0.0806 (0.0766,0.0847)	0.0736 (-0.1044,0.2516)	0.0701 (-0.0805,0.2209) 81.8%
South Atlantic	0.0653 (0.0633,0.0673)	0.0578 (0.0558,0.0598)	0.0549 (-0.0123,0.1253)	0.0558 (-0.0092,0.1227) 95.5%
West North Central	0.0627 (0.06,0.0653)	0.1483 (0.1449,0.1518)	0.1243 (-0.0519,0.3062)	0.1079 (-0.0419,0.2596) 92.1%
West South Central	0.0597 (0.0575,0.062)	0.0812 (0.0789,0.0835)	0.0798 (0.0366,0.1225)	0.08 (0.0389,0.1214) 99.99%

Notes:

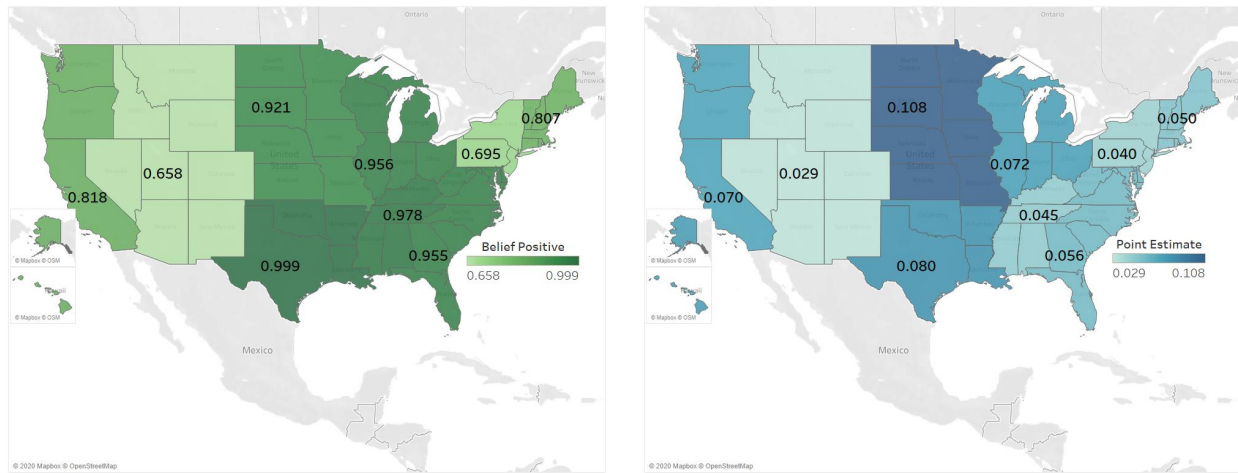
[1] These are point estimates and 95% confidence intervals obtained by using standard techniques. All results are "significant". The OLS results indicate that there is a positive empirical relationship between education and earnings, although this relationship cannot be interpreted as causal. The positive point estimates for the IV estimates indicates that there might be a positive causal relationship between education and earnings. A standard interpretation of the 95% confidence interval as derived from the asymptotic approximation of the distribution of the IV-2SLS estimate would lead us to believe that the returns to education are "significantly" positive for all Census Regions.

[2] This mean and 95% credible interval are gathered by Gibbs sampling 10,000 times from the posterior of the instrumental variables estimate. We use a "weak" prior that assumes the coefficients have variance 1 million (the ridge regression penalty is 1/(1 million)). The point estimates are the mean of the Gibbs samples, while the 95% credible intervals have endpoints determined by the 2.5 and 97.5 percentiles of the samples. Note that the point estimates for the case with weak priors are quite similar to the "standard" IV point estimates, but the 95% credible intervals are so much wider that only one out of nine intervals doesn't include zero. The wider intervals lead us to conclude that it is inappropriate to assume that the "standard" 2SLS estimate is normally distributed with the variance as derived from the delta method.

[3] As in [2], this mean and 95% credible interval are gathered by Gibbs sampling 10,000 times from the posterior of the instrumental variables estimate. We use K-Fold cross validation to determine that a prior variance of 4.45 for the second stage and upwards of 1 million for the first stage performed the best at out-of-sample prediction (minimizing mean squared error using a two-stage approach).

[4] The values in the last column display the percent of the posterior that lies above zero. This is a measure of our confidence that the returns to education are positive. We conclude that our identification strategy leads us to be 95% confident that the returns to education are positive for the following four census regions: East North Central, East South Central, South Atlantic, West South Central. Note that our largest point estimate (West North Central) is not 95% credible!

Figure 4: The Heterogenous Returns to Education



(a) Percent of Posterior Distribution Positive

(b) Mean of Posterior Distribution

These two maps show summary measures of the posterior distribution of the causal returns to education. We use K-Fold cross validation to determine that a prior variance of 4.5 for the second stage and a large prior variance upwards of 1000 for the first stage performed the best at out-of-sample prediction (using a two-stage approach) for log weekly earnings.

Figure 4a shows the proportion of the posterior that lies above zero, representing the posterior chance that the returns to education are positive. Our identification strategy leads us to be 95% confident that the returns to education are positive for the following four census regions: East North Central, East South Central, South Atlantic, West South Central.

Figure 4b shows the mean of the posterior distribution. Note that a larger mean does not correspond to more confidence in positive returns to education. Although West North Central had the largest point estimate, we aren't even 95% sure that the returns to education are positive for West North Central.

7 Conclusion

This article estimates the heterogeneity in the returns to education. We follow Angrist and Keueger (1991) in using the quarter-of-birth as an instrument for the level of schooling obtained by men born between 1930 and 1940. The quarter-of-birth is a "weak instrument" in the sense that we cannot assume the IV-2SLS estimate is distributed normal. This assumption of normality is an appeal to the asymptotic distribution of the two-stage least squares estimate, and calculations of the "standard errors" that are displayed in common analytical programs like Stata report values calculated from applying the Delta-method.

Instead of using the asymptotic distribution, we sample from the posterior distribution of the two-stage least squares estimate. When using a weak prior, our point estimates agree with the two-stage least squares instrumental variables estimates. However our "95%" confidence intervals are much wider, leading us to reject positive returns to school-

ing for certain regions, while the standard approach would lead to strong, "significant" results for all the regions. We also use cross-validation to determine the priors on the coefficients (same as ridge-regression regularization penalties) that perform the best for out-of-sample prediction using a two-stage least squares approach. The cross validation tells us to use an weak prior for the first stage, but use a stronger prior in the second stage. This empirical result agrees with the observation in [Newey and Powell \(2003\)](#) that some regularization is needed in the second stage to solve the ill-posed inverse problem.

Our specific application is complicated by the necessity of estimating heterogeneous treatment effects. In this article, we interact our proposed treatment (education level) with geographic characteristics. We show that the currently accepted method of only have one first stage equation - which instruments for education - can lead to biased results if the source of heterogeneity is correlated with the omitted factors in the first stage specification. In our case, instrumental variables estimate would still be biased if the region of birth is correlated with any omitted factor that determined the level of schooling. Since this assumption is too strong, instead we instrument for all the heterogeneous treatments, leading us to have multiple first stage estimating equations. Instead of deriving the actual posterior distribution, we iteratively sample from the conditional posterior distributions.

We have two main findings: for men born between 1930 and 1940 and who were induced to attend an additional year of schooling due to the the interaction between compulsory schooling laws and their quarter of birth, we are 95% sure that the returns to education were positive for these four regions, and the returns to education were probably highest for this region ⁴.

⁴Please refer to [Figure 2](#) for a visual display of the regions, or [Table 2](#) for the exact definitions.

References

- Andrews, I., J. H. Stock, and L. Sun (2019). Weak instruments in instrumental variables regression: Theory and practice. *Annual Review of Economics* 11, 727–753.
- Angrist, J. D. and A. B. Keueger (1991). Does compulsory school attendance affect schooling and earnings? *The Quarterly Journal of Economics* 106(4), 979–1014.
- Bekker, P. A. (1994). Alternative approximations to the distributions of instrumental variable estimators. *Econometrica: Journal of the Econometric Society*, 657–681.
- Bettinger, E., O. Gurantz, L. Kawano, B. Sacerdote, and M. Stevens (2019, February). The long-run impacts of financial aid: Evidence from california’s cal grant. *American Economic Journal: Economic Policy* 11(1), 64–94.
- Bound, J., D. A. Jaeger, and R. M. Baker (1995). Problems with instrumental variables estimation when the correlation between the instruments and the endogenous explanatory variable is weak. *Journal of the American statistical association* 90(430), 443–450.
- Card, D. (1995). Using geographic variation in college proximity to estimate the return to schooling, aspects of labour market behaviour: essays in honour of john vanderkamp. ed. *LN Christofides, EK Grant, and R. Swidinsky*.
- Card, D. (1999). The causal effect of education on earnings. 3, *Part A*, 1801–1863.
- Chamberlain, G. and G. Imbens (1996). Hierarchical bayes models with many instrumental variables. *NBER Technical Working Paper*.
- Chao, J. C. and P. C. Phillips (1998). Posterior distributions in limited information analysis of the simultaneous equations model using the jeffreys prior. *Journal of Econometrics* 87(1), 49–86.
- Dreze, J. H. (1976). Bayesian limited information analysis of the simultaneous equations model. *Econometrica: Journal of the Econometric Society*, 1045–1075.
- EconML, M.-R. (2019). EconML: A Python Package for ML-Based Heterogeneous Treatment Effects Estimation. <https://github.com/microsoft/EconML>. Version 0.x.
- Fieller, E. C. (1932). The distribution of the index in a normal bivariate population. *Biometrika*, 428–440.

- Hahn, J. and J. Hausman (2002). A new specification test for the validity of instrumental variables. *Econometrica* 70(1), 163–189.
- Hartford, J., G. Lewis, K. Leyton-Brown, and M. Taddy (2017). Deep iv: A flexible approach for counterfactual prediction. In *Proceedings of the 34th International Conference on Machine Learning-Volume 70*, pp. 1414–1423. JMLR. org.
- Hoogerheide, L., F. Kleibergen, and H. K. van Dijk (2007). Natural conjugate priors for the instrumental variables regression model applied to the angrist–krueger data. *Journal of Econometrics* 138(1), 63–103.
- Leamer, E. E. (2010). Tantalus on the road to asymptopia. *Journal of Economic Perspectives* 24(2), 31–46.
- Link, W. A. and M. J. Eaton (2012). On thinning of chains in mcmc. *Methods in ecology and evolution* 3(1), 112–115.
- Lopes, H. F. and N. G. Polson (2014). Bayesian instrumental variables: priors and likelihoods. *Econometric Reviews* 33(1-4), 100–121.
- MacEachern, S. N. and L. M. Berliner (1994). Subsampling the gibbs sampler. *The American Statistician* 48(3), 188–190.
- Nelson, C. R. and R. Startz (1990). The distribution of the instrumental variables estimator and its t-ratio when the instrument is a poor one. *Journal of business*, S125–S140.
- Newey, W. K. and J. L. Powell (2003). Instrumental variable estimation of nonparametric models. *Econometrica* 71(5), 1565–1578.
- Phillips, P. C. (2009). Exact distribution theory in structural estimation with an identity. *Econometric Theory* 25(4), 958–984.
- Sanderson, E. and F. Windmeijer (2016). A weak instrument f-test in linear iv models with multiple endogenous variables. *Journal of Econometrics* 190(2), 212–221.
- Staiger, D. and J. H. Stock (1994). Instrumental variables regression with weak instruments. Technical report, National Bureau of Economic Research.
- Stock, J. H. and M. Yogo (2002). Testing for weak instruments in linear iv regression. Technical report, National Bureau of Economic Research.

- Wiesenfarth, M., C. M. Hisgen, T. Kneib, and C. Cadarso-Suarez (2014). Bayesian non-parametric instrumental variables regression based on penalized splines and dirichlet process mixtures. *Journal of Business & Economic Statistics* 32(3), 468–482.
- Young, A. (2019). Consistency without inference: Instrumental variables in practical application.
- Zellner, A. (1978). Estimation of functions of population means and regression coefficients including structural coefficients: A minimum expected loss (melo) approach. *Journal of Econometrics* 8(2), 127–158.
- Zellner, A., T. Ando, N. Baştürk, L. Hoogerheide, and H. K. Van Dijk (2014). Bayesian analysis of instrumental variable models: Acceptance-rejection within direct monte carlo. *Econometric Reviews* 33(1-4), 3–35.