

Appendix

Parameters of Literature Review

We searched every issue of the APSR and AJPS on Google Scholar between 2008 and 2015 for the term “fixed effects,” which returned roughly 290 articles. We trimmed this sample to include only observational studies (no randomized experiments, in which fixed effects are used merely for efficiency gains and are orthogonal to treatment assignment by design) that used the term “fixed effects” in the traditional econometric sense, i.e. mutually exclusive and exhaustive group indicators such as units, time periods, regions, etc., (as opposed to a non-varying coefficient in a hierarchical model). We retained studies which used fixed effects in a core analysis of a central hypothesis based on our reading of the manuscript (e.g. not simply in a footnoted robustness check). We also excluded methods, theoretical and review articles and all articles that did not use plain ordinary least squares regression. Given our interest is in how researchers formulate counterfactuals, we also excluded articles which featured a dichotomous treatment, since there is little discretion for counterfactuals in the case of a 0/1 treatment (i.e. the treatment either occurs or it does not). The end result was a sample of 55 articles. We then read the sections of each study outlining the research design and empirical results and coded each study for whether (1) the article explicitly discussed the variation being used during estimation (i.e. “within countries over time”...); (2) the article discussed a counterfactual when discussing results; and (3) the counterfactual shift that was considered in X was chosen based on the transformed distribution of X (e.g. the within-unit distribution in the case of unit fixed effects). We counted articles which explicitly mentioned the within-unit variation being used, or that described fixed effects as “controlling for between-unit differences,” as satisfying the first criteria, which we regard as lenient. While we cannot guarantee that we did not miss a relevant case, we think this sample provides a good portrait of the use and interpretation of fixed effects estimators in the top tier of the field today.

Counterfactuals from Published Examples

Below are the counterfactuals discussed in the published examples we highlight in Table 1. Note that it is not uncommon for several counterfactuals to be discussed in one article, but we discuss only one from each study listed in Table 1.

Berrebi and Klor 2008; a one-unit shift in X

“We find that the occurrence of a terror fatality within three months of the elections is associated with a .45 percentage point increase in the locality’s relative electoral support for the right bloc of political parties, ” (287).

Gasper and Reeves; a one-unit shift in X

“Column 3 of Table 1 shows that a turn down is worth 2.7 points for an incumbent governor’s vote share,” (350).

Ichino and Nathan 2009; a one standard deviation shift in X

“...a one standard deviation increase in the spatially weighted population share of Akans beyond a polling station (about 0.21) results in a predicted 6.9 percentage point greater NPP presidential vote share (95% CI: [5.7, 10.6]),” (351-352).

Scheve and Stasavage 2012; a one-unit/min. to max. shift in X

“... a country which mobilized for mass warfare for an entire five-year period increased its top inheritance tax rate by 18 to 23 percentage points compared to a country that did not mobilize for war,” (94).

Snyder and Strömberg 2010; a one-unit/min. to max. shift in X

“... a change from the lowest to the highest values of congruence is associated with a 28 percent increase in the probability of correctly recalling a candidate’s name. This is about as large as the effect of changing a respondent’s education from grade school to some college,” (372).

Appendix: Fixed Effects Analysis in R

Jonathan Mummolo and Erik Peterson

The following code implements our recommendations for analyzing and interpreting fixed effects regression models. We will use Snyder and Strömberg (2010), discussed in the main text, as a working example.

1. Isolate Relevant Variation

First set your working directory and load the data.

```
setwd("~/Dropbox/FE Paper/Replication Data/Snyder and Stromberg 2010")
load("congre recall.Rdata")
## rename dataframe
d <- congre recall
```

The following code replicates the second model in Table 4, Column 1 in Snyder and Strömberg (2010). The outcome is an indicator for whether an individual recalled their congressional representative's name, and the treatment, "cov2c," is media market congruence (range 0 to 1, continuous). The reported coefficient on the treatment in the published article is 0.28.

```
## estimate model of Name Recall with year and
## incumbent fixed effects
reg <- lm(cong_recall ~ cov2c + factor(year) + factor(incnum),
  data = d)
## view first two coefficients
summary(reg)$coefficients[1:2, ]
```

```
##           Estimate Std. Error t value Pr(>|t|)
## (Intercept) 0.03346182 0.10520300 0.318069 7.504376e-01
## cov2c       0.27878620 0.05459314 5.106616 3.325184e-07
```

```
## store estimation data
est <- reg$model
## make an object that is the variable name of the
## treatment
x <- "cov2c"
## store the coefficient estimate
b <- reg$coefficients[x]
```

We have replicated the result. Let's isolate the variation in the treatment being used to estimate this coefficient by residualizing the treatment with respect to the fixed effects.

```
x.resid <- lm(cov2c ~ factor(year) + factor(incnum),
  data = d)$residuals
```

If this is the same variation being used to estimate the coefficient of interest, then we should be able to regress the outcome on this transformed treatment and get the same coefficient.

```
reg2 <- lm(est$cong_recall ~ x.resid)
summary(reg2)$coefficients[1:2, ]
```

```
##              Estimate Std. Error  t value    Pr(>|t|)
## (Intercept) 0.3139543 0.003900063 80.499813 0.000000e+00
## x.resid      0.2787862 0.057599513  4.840079 1.311586e-06
```

We indeed get the same result. Note that we did not residualize the outcome, but if we did, we would also get the same coefficient on the treatment.

```
y.resid <- lm(cong_recall ~ factor(year) + factor(incnum),
              data = d)$residuals
reg3 <- lm(y.resid ~ x.resid)
summary(reg3)$coefficients[1:2, ]
```

```
##              Estimate Std. Error  t value    Pr(>|t|)
## (Intercept) 5.870861e-17 0.003583651 1.638235e-14 1.000000e+00
## x.resid      2.787862e-01 0.052926467 5.267425e+00 1.403801e-07
```

2. Identify a Plausible Shift in X

Now that we have isolated the relevant variation in the treatment and estimated the coefficient of interest, we need to determine what number to multiply the coefficient by to judge whether the treatment effect is substantively significant. The coefficient itself represents the effect of a one-unit shift in the treatment, which in this case is a shift from the theoretical minimum to the theoretical maximum of the variable. Is this a plausible shift?

```
range(est[, x])
```

```
## [1] 0.01890017 0.96752632
```

```
sdx <- sd(est[, x])
sdx
```

```
## [1] 0.2860724
```

```
sdxresid <- sd(x.resid)
sdxresid
```

```
## [1] 0.06771239
```

```
sdx/sdxresid
```

```
## [1] 4.224816
```

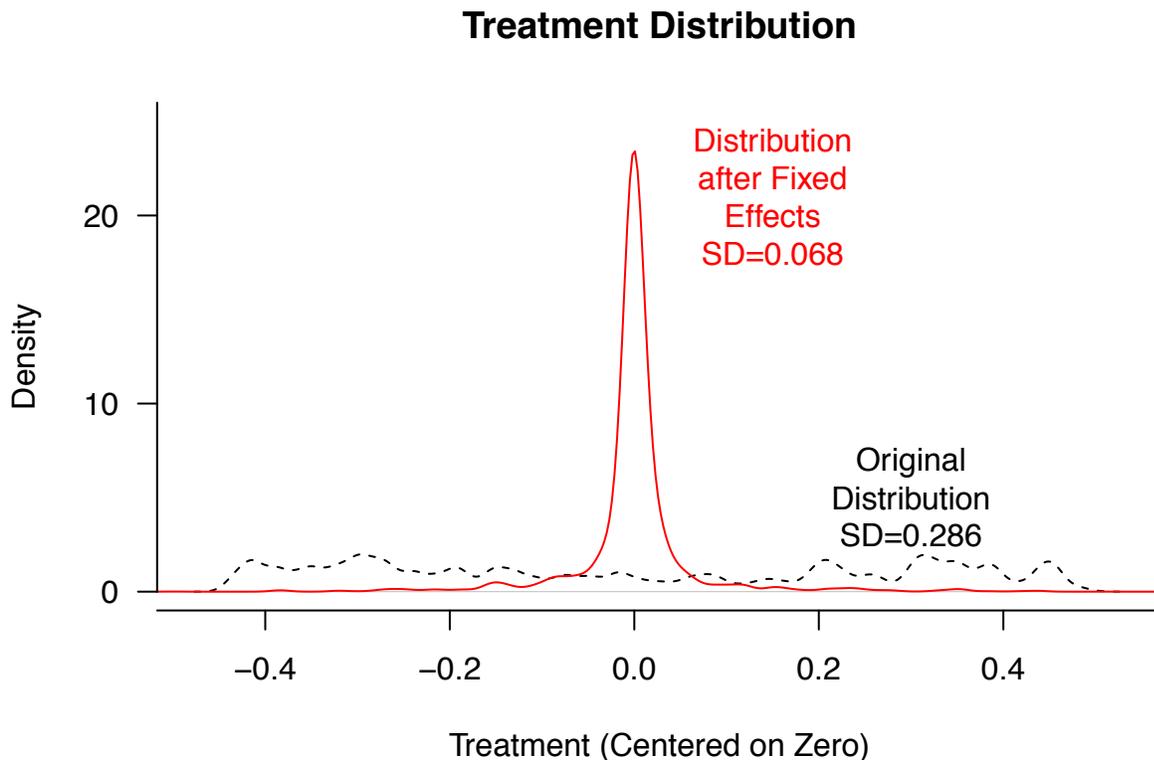
```
## compute % difference
((sdxresid - sdx)/sdx) * 100
```

```
## [1] -76.33033
```

From this basic analysis we learn that a 0-to-1 shift is beyond the observed range of the data [.019, .968]. Further, while the standard deviation of the original treatment is 0.29, the standard deviation of the transformed treatment used during estimation is 0.068, a reduction of 76%.

Let's visualize this reduction in variance. Overlapping density plots are useful for this. We will also center both variables on zero by demeaning them to make for an easier comparison of their spreads.

```
plot(density(x = est[, x] - mean(est[, x]), bw = 0.01),
     col = "black", cex = 0.3, lty = 2, main = "Treatment Distribution",
     xlab = "Treatment (Centered on Zero)", ylim = c(0,
     25), axes = F)
lines(density(x = x.resid - mean(x.resid), bw = 0.01),
     col = "red", cex = 0.7, xlab = "")
text(x = 0.15, y = 22, labels = "Distribution\nafter Fixed\nEffects",
     col = "red")
text(x = 0.15, y = 18, labels = paste("SD=", round(sdxresid,
     3), sep = ""), cex = 1, col = "red")
text(x = 0.3, y = 6, labels = "Original\nDistribution",
     col = "black")
text(x = 0.3, y = 3, labels = paste("SD=", round(sdx,
     3), sep = ""), cex = 1)
axis(1, seq(-1, 1, by = 0.2))
axis(2, seq(0, 40, by = 10), las = 2)
```



The plot makes very clear that a drastic reduction in variation has occurred as a result of the fixed effects being employed.

It is also useful to examine the distribution of within-unit ranges when trying to identify plausible counterfactuals. Let's generate that distribution and plot it as well. First, we will residualize the data with respect to

year dummies. The variation that remains is the within-incumbent variation used during estimation.

```
bounds <- list(NA)
ranges <- NA
units <- unique(est[, "factor(incnum)"])
est$x.resid2 <- lm(est[, x] ~ est[, "factor(year)"])$residuals

for (i in 1:length(units)) {

  bounds[[i]] <- range(est[est[, "factor(incnum)"] ==
    units[i], "x.resid2"])
  ranges[i] <- bounds[[i]][2] - bounds[[i]][1]

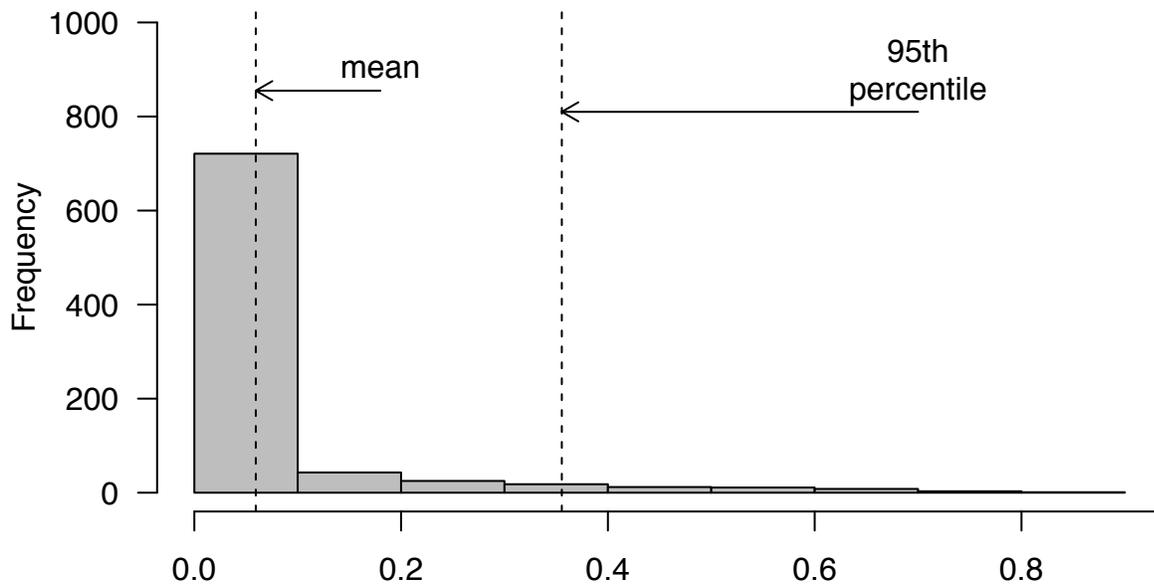
}

## What percent of the within-unit ranges are zero?
(length(ranges[ranges == 0])/length(ranges)) * 100
```

```
## [1] 44.18052
```

```
hist(ranges, col = "grey", main = "Within-Unit Ranges\nof Treatment",
  xlab = "Within-Incumbent Ranges", ylim = c(0, 1000),
  axes = F)
axis(1, seq(0, 1, by = 0.2))
axis(2, seq(0, 1000, by = 200), las = 2)
abline(v = mean(ranges), lty = 2)
abline(v = quantile(ranges, probs = 0.95), lty = 2)
text(x = 0.18, y = 900, labels = "mean")
text(x = 0.7, y = 900, labels = "95th\npercentile")
arrows(x0 = 0.18, x1 = mean(ranges), y0 = 855, y1 = 855,
  length = 0.1)
arrows(x0 = 0.7, x1 = quantile(ranges, probs = 0.95),
  y0 = 810, y1 = 810, length = 0.1)
```

Within-Unit Ranges of Treatment



Within-Incumbent Ranges

As the histogram shows, the range of the treatment in the bulk of units is concentrated near zero, with a mean of roughly 0.06 and a 95th percentile of about 0.36.

With the above quantities in hand, researchers can now decide which counterfactual shift in the treatment they would like to discuss. They can also provide a sense of the plausibility of that counterfactual shift by describing where in the within-unit distribution the chosen shift is located.

```
## A within-unit SD shift  
b * sdx
```

```
##      cov2c  
## 0.07975304
```

```
## A shift the size of the mean within-unit range  
b * mean(ranges)
```

```
##      cov2c  
## 0.01657017
```

```
## A shift the size of the 95th percentile in  
## within-unit ranges  
b * quantile(ranges, probs = 0.95)
```

```
##      cov2c  
## 0.09908938
```

```
## A shift the size of the maximum within-unit range
b * max(ranges)
```

```
##      cov2c
## 0.2458807
```

In addition, we learn that 44% of the within-unit ranges are zero. These units do not contribute to the fixed effects estimation. (That is, if we dropped these units from the data, we would get the same coefficient estimate on the treatment.) Researchers should note the share of units actually contributing to the coefficient estimate when describing results, as this informs the population being studied and the generalizability of results.

```
no.var <- which(ranges == 0)
## Drop non-varying units from the data
est2 <- subset(est, !(est[, "factor(incnum)"] %in%
  no.var))
## Re-estimate the model
reg.drop <- lm(cong_recall ~ cov2c + est2[, "factor(incnum)"] +
  est2[, "factor(year)"], data = est2)
## We get the same coefficient:
summary(reg.drop)$coefficients[1:2, ]
```

```
##              Estimate Std. Error t value    Pr(>|t|)
## (Intercept) 0.03346182 0.10520300 0.318069 7.504376e-01
## cov2c       0.27878620 0.05459314 5.106616 3.325184e-07
```

References

Snyder, James M. and David Strömberg. 2010. "Press Coverage and Political Accountability." *Journal of Political Economy*, 118(2), 355-408.