

The Fox News Effect

5 August 2015

Recently, many scholars have been interested in quantifying the effect of the national news media on the behavior of electorate. Understanding how the national media environment affects the election results is critically important, but also notoriously difficult. In particular, even if an association between media coverage and election outcomes exists, it is difficult to identify whether this association is due to the media's influence over the voters or the result of the media adjusting its contents to the preferences of voters.

In this exercise, we will consider the entry of the Fox News Channel (hereafter Fox News) into the television market in the late 1990s. This exercise is based on the following study:

Stefano DellaVigna and Ethan Kaplan (2007). "The Fox News Effect: Media Bias and Voting." *Quarterly Journal of Economics*, 122:3, pp.1187-1234.

Note that due to the nature of negotiations between cable companies and television networks, adding a new channel to the line-up of a cable company may take a long time. For this reason, in contrast to what many people might expect, the Fox News was not able to enter conservative media markets first. We will be looking at some of the differences between the towns that initially did not receive Fox News, and those that did. The data set is in the csv file `foxnews.csv`. It contains information for 10,126 towns across 28 states in the United States:

Name	Description
<code>town</code>	Town name
<code>state</code>	State in which the town is located
<code>subrf2000</code>	Share of Fox News subscribers in 2000
<code>gopvoteshare2000</code>	Two-party vote share for the Republicans (2000 Presidential election)
<code>gopvoteshare1996</code>	Two-party vote share for the Republicans (1996 Presidential election)
<code>gopvoteshare1992</code>	Two-party vote share for the Republicans (1992 Presidential election)
<code>college1990</code>	Proportion of population with a college degree in 1990
<code>male1990</code>	Proportion of male population in 1990
<code>black1990</code>	Proportion of black population in 1990
<code>hisp1990</code>	Proportion of hispanic population in 1990
<code>income1990</code>	Median income in 1990
<code>logincome1990</code>	Median income in 1990 on the logarithmic scale

Question 1

We will investigate whether there are any systematic differences in the distribution of some key pre-treatment variables (`hs1990`, `black1990`, `hisp1990`, `male1990`, `logincome1990`) between towns that received Fox News as compared to those that did not. First, create a new variable called `foxnews2000` that takes the value of 1 if the share of Fox News subscribers in a given town is strictly larger than 0 and equals 0 otherwise. Create five Quantile-Quantile plots to assess the similarity of the distributions for towns with and without subscribers across the five variables (`hs1990`, `black1990`, `hisp1990`, `male1990`, `logincome1990`). Interpret the results. Are there any consistent patterns of differences between the two groups of towns? What do the plots tell you about our ability to make causal inferences regarding the effect of Fox News on the election outcome?

Answer 1


```
par(cex = 1.5)
## Read the data
foxnews <- read.csv("data/foxnews.csv")

## Fox News Indicator
foxnews$foxnews2000 <- ifelse(foxnews$subrf > 0, 1, 0)

## Proportion of College graduates in 1990
qqplot(foxnews$college1990[foxnews$foxnews2000 == 1],
       foxnews$college1990[foxnews$foxnews2000 == 0],
       xlab = "Towns with access to Fox News",
       ylab = "Towns with no access to Fox News",
       main = "Proportion of College graduates in 1990")
abline(0,1, col = "red")
```

Towns with no access to Fox News

Proportion of College graduates in 1990

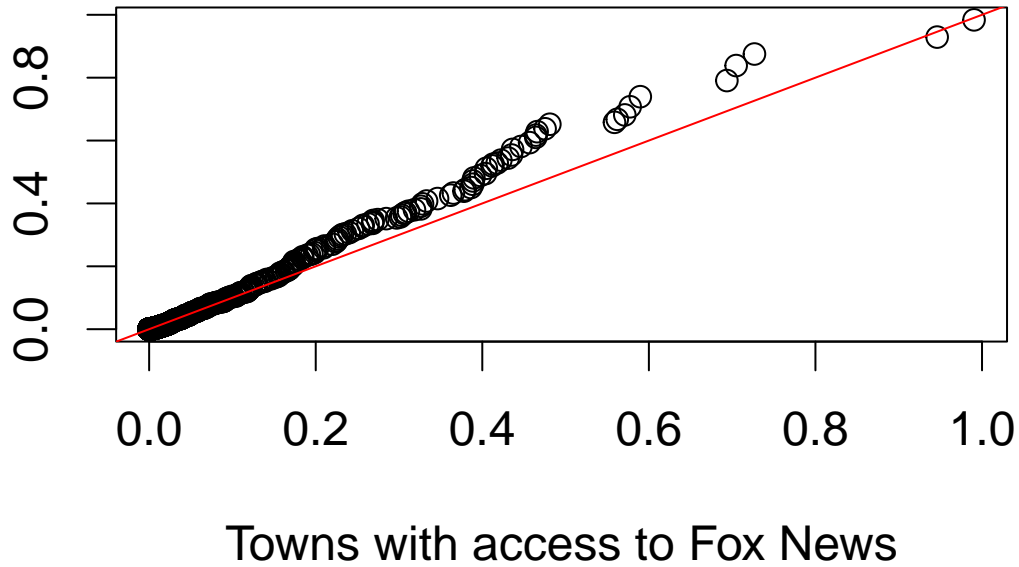


Towns with access to Fox News

```
## Proportion of Blacks in 1990
qqplot(foxnews$black1990[foxnews$foxnews2000 == 1],
       foxnews$black1990[foxnews$foxnews2000 == 0],
       xlab = "Towns with access to Fox News",
       ylab = "Towns with no access to Fox News",
       main = "Proportion of Blacks in 1990")
abline(0,1, col = "red")
```

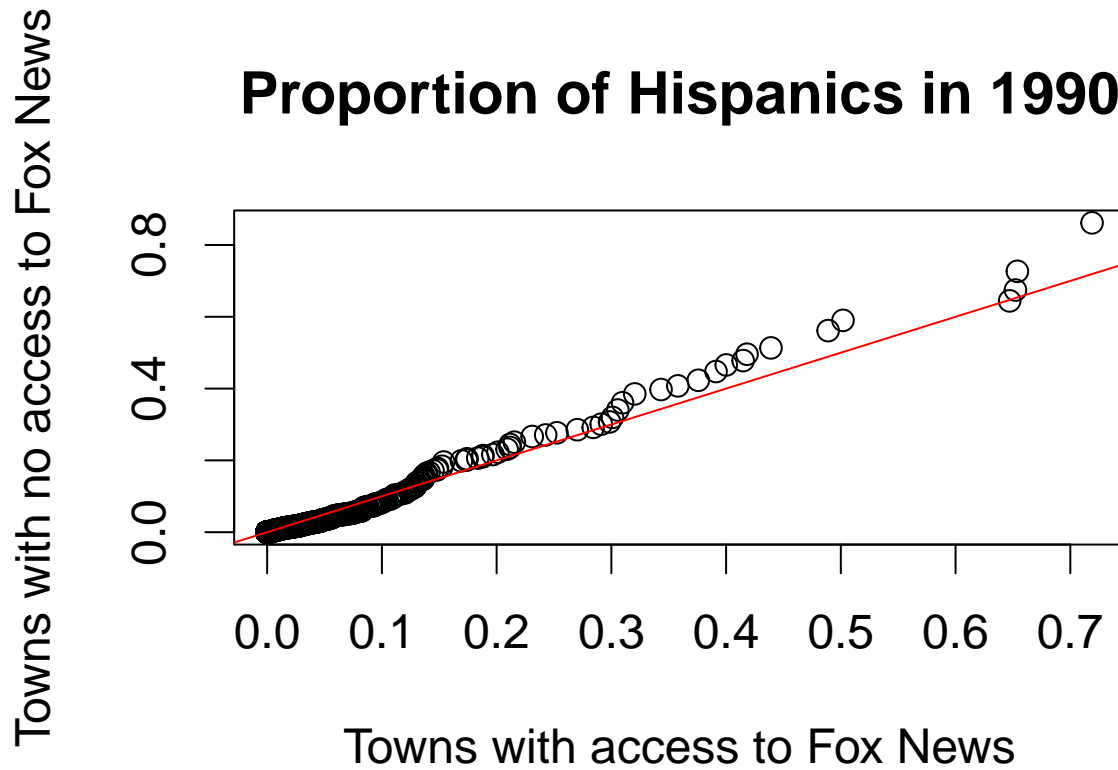
Towns with no access to Fox News

Proportion of Blacks in 1990



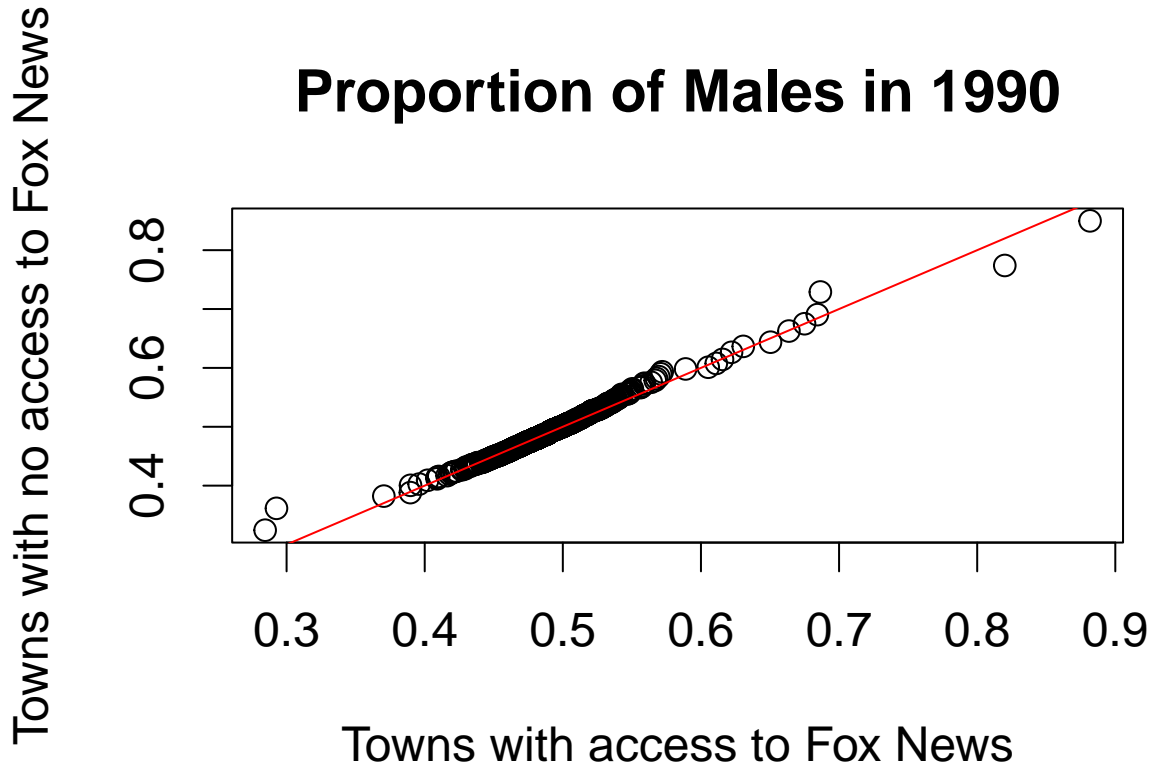
```
## Proportion of Hispanics in 1990
qqplot(foxnews$hispanic1990[foxnews$foxnews2000 == 1],
       foxnews$hispanic1990[foxnews$foxnews2000 == 0],
       xlab = "Towns with access to Fox News",
       ylab = "Towns with no access to Fox News",
       main = "Proportion of Hispanics in 1990")
abline(0,1, col = "red")
```

Proportion of Hispanics in 1990



```
## Proportion of Males in 1990
qqplot(foxnews$male1990[foxnews$foxnews2000 == 1],
foxnews$male1990[foxnews$foxnews2000 == 0],
       xlab = "Towns with access to Fox News",
       ylab = "Towns with no access to Fox News",
       main = "Proportion of Males in 1990")
abline(0,1, col = "red")
```

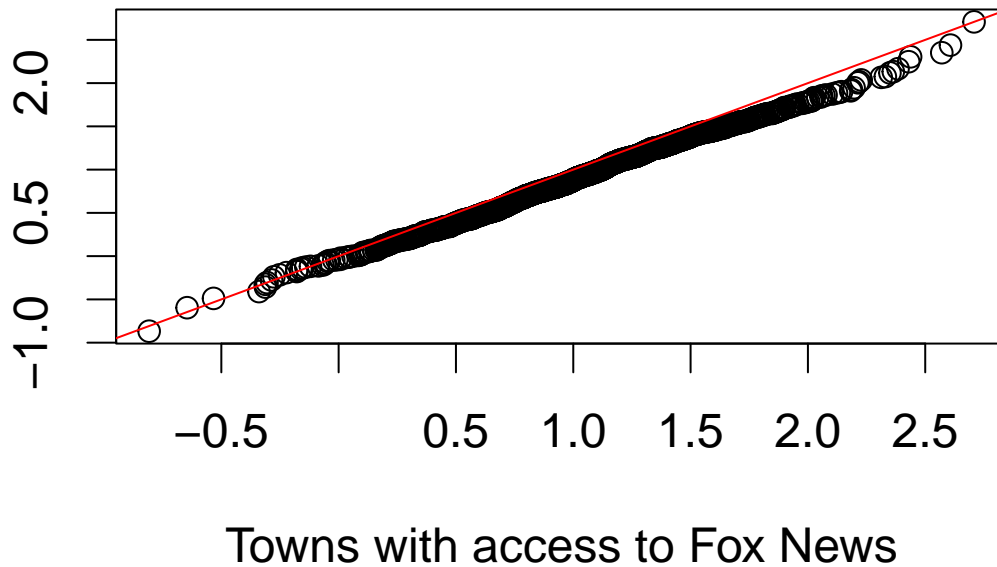
Proportion of Males in 1990



```
## Logarithm of the Median Income in 1990
qqplot(foxnews$logincome1990[foxnews$foxnews2000 == 1],
       foxnews$logincome1990[foxnews$foxnews2000 == 0],
       xlab = "Towns with access to Fox News",
       ylab = "Towns with no access to Fox News",
       main = "Logarithm of the Median Income in 1990")
abline(0,1, col = "red")
```

Towns with no access to Fox News

Logarithm of the Median Income in 1990



The distributions for the proportions of high-school graduates and proportion of males are nearly identical across the towns that received the Fox News signal as compared to those that did not. The identity is not perfect: for example, the lower quantiles for the proportion of high school graduates are more skewed toward higher values for towns with no access to Fox News, while the reverse is true for the upper quantiles. There is more imbalance along the other three covariates. For the logarithm of the median income, the quantiles lie consistently below the 45-degree line, indicating that the income distribution for towns that received the Fox News signal is systematically shifted to higher values as compared to those that did not. Similar, though less pronounced, imbalance also obtains for the proportions of Hispanic and Black residents. These differences between the distributions confirm that exposure to Fox News cannot be considered entirely random, but the feasibility of the as-if random assumption can be defended due to the small magnitude of the imbalance on the observed covariates. These differences imply that any estimate of the causal effect of exposure to Fox News should, as a further check on our analysis, account for any differences in pre-treatment values on the outcome variable across towns that received the Fox News signal and those that did not. The difference-in-differences design would be an appropriate choice for that, subject to the validity of the parallel trend assumption.

Question 2

We further examine whether there are any clear differences between those towns that did receive Fox News and those that did not. To do this, apply the k -means algorithm with two clusters to the five variables you analyzed in the previous question. Be sure to remove any missing values and scale each variable so that their means are zero and standard deviations are one, before applying the algorithm. What is the distribution of the clusters with respect to whether or not towns received Fox News? What are the characteristics of each cluster? Explain how this analysis answers the question about our ability to make causal inferences about the electoral effect of Fox News.

Answer 2

```
par(cex = 1.5)
## Setting a seed for replicability
```

```

set.seed(12345)
## Subset the data
foxnews.subset <- na.omit(foxnews[,c("foxnews2000", "male1990",
                                   "hispanic1990", "logincome1990",
                                   "black1990", "hispanic1990")])

## Scale the data
foxnews.subset.scaled <- scale(foxnews.subset)
## k-means
cluster.2 <- kmeans(foxnews.subset.scaled[,2:6], centers = 2)
## Means for each variable by cluster
cluster.2$centers

##      male1990  hispanic1990 logincome1990  black1990 hispanic1990.1
## 1 -0.008833611 -0.1226497   0.003798685 -0.00464802 -0.1226497
## 2  0.409162629  5.6809936  -0.175950696  0.21529090  5.6809936

## Means on the original scale
colMeans(foxnews.subset[cluster.2$cluster == 1, 2:6])

##      male1990      hispanic1990 logincome1990      black1990      hispanic1990.1
## 0.49065584      0.01258006      0.93806252      0.03121454      0.01258006

colMeans(foxnews.subset[cluster.2$cluster == 2, 2:6])

##      male1990      hispanic1990 logincome1990      black1990      hispanic1990.1
## 0.50269360      0.34992329      0.85998481      0.05226992      0.34992329

## Cross Table Fox vs Clusters
fox.clusters <- table(cluster.2$cluster,
                      foxnews.subset$foxnews2000)
fox.clusters

##
##      0      1
## 1 7792 1796
## 2  173   34

## The rate of exposure to Fox News is similar across clusters
fox.clusters[1,2]/sum(fox.clusters[1,])

## [1] 0.1873175

fox.clusters[2,2]/sum(fox.clusters[2,])

## [1] 0.1642512

```

It can be noted that if compared to cluster 2, towns in cluster 1 have around the same number of males, less hispanics and blacks, while their income is larger. However, in terms of the presence of Fox News across these two clusters, we can see it is quite similar in relative terms i.e., around 18 percent of the towns in each cluster received the Fox News signal, while the rest did not.

Question 3

We begin to examine the relationship between the exposure to Fox News in 2000 and the change in the GOP's vote share from the 1996 to the 2000 Presidential election. First, create a new variable that measures the difference between the Republican vote share in 2000 and in 1996. Compute the correlation between this new variable and `subrf2000` and provide an interpretation of the result.

Answer 3

```
## Difference in the vote shares 2000 - 1996
foxnews$voteshare.diff <- foxnews$gopvoteshare2000 -
                        foxnews$gopvoteshare1996

## Correlation:
corr <- cor(foxnews$voteshare.diff, foxnews$subrf2000,
            use = "complete.obs")
corr
```

```
## [1] -0.1066677
```

The correlation between the two variables stands at -0.107, which suggests a weak negative association. This means that greater exposure to Fox News corresponds to (but does not necessarily bring about!) a lower electoral gain for the Republicans. The scatterplot helps to visualize this conclusion.

Question 4

We now estimate the causal effect of Fox News on the Republicans' vote share. For this question, use `foxnews2000` to measure exposure to Fox News. Interpret the results. What estimation strategy did you use to identify this causal effect? What is the assumption required for this analysis to be valid? Interpret this assumption in the context of this particular question. In your view, how credible is this assumption? Use the 1992 and 1996 election outcomes, both of which took place before the creation of the Fox News channel, to probe the credibility of the assumption.

Answer 4

```
## Difference-in-Differences
DiD.treat <- mean(foxnews$voteshare.diff[foxnews$foxnews2000 == 1],
                 na.rm = TRUE)
DiD.control <- mean(foxnews$voteshare.diff[foxnews$foxnews2000 == 0],
                  na.rm = TRUE)

DiD <- DiD.treat - DiD.control
DiD
```

```
## [1] -0.01298383
```

```
## Placebo Test
foxnews$voteshare.diff.2 <- foxnews$gopvoteshare1996 -
                        foxnews$gopvoteshare1992

Placebo.t <- mean(foxnews$voteshare.diff.2[foxnews$foxnews2000 == 1],
                 na.rm = TRUE)
Placebo.c <- mean(foxnews$voteshare.diff.2[foxnews$foxnews2000 == 0],
                 na.rm = TRUE)

Placebo <- Placebo.t - Placebo.c
Placebo
```

```
## [1] 0.0071524
```

According to our difference-in-differences estimate the entry of Fox News made things worse for the Republicans as it reduced the vote share in favor of the GOP by more than 1 percentage point.

The main assumption behind difference-in-differences is the 'parallel paths' assumption. Basically, what we

need to assume is that in the absence of the treatment, any pre-existing differences in the outcome of interest would have remained constant across treatment and control groups. The latter is key to identify the effect of Fox News, as any movement beyond the pre-existing differences in towns that did and did not received the signal can be attributed to the entry of Fox News. Given that we could not find any major difference in other relevant covariates between towns that received the Fox News signal and towns that did not, it seems reasonable to make this assumption.

To assess the validity of this assumption, we perform a ‘placebo test.’ We do so to discard the possibility that the effect we found is not due to an unobserved trend in the electorate agaiings the Republicans. In other words, Fox News should not have an impact before its entry to the cable market. If we do find an effect, then there might be something different in the towns we are comparing. As it can be noted above, the placebo test reveals no effect on difference in the vote shares for the Republicans from 1992 to 1996, when comparing towns that received the signal of Fox News in the year 2000 with towns that did not.

Question 5

We further divide the towns that received Fox News into three groups based on the share of Fox News subscribers. Among the towns who received Fox News, the ‘High exposure’ group represents the group of towns whose share of subscribers is greater than or equal to the 66 percentile (among those who received Fox News). In contrast, the ‘Low exposure’ group represents the group of towns whose share of subscribers is less than or equal to the 33 percentile (among those who received Fox News). Conduct the same analysis as in the previous question but separately for the ‘High exposure’ and ‘Low exposure’ groups where the control group is the ‘No exposure’ group. Interpret the results.

Answer 5

```
## 33rd Percentile
p33 <- quantile( foxnews$subrf[ foxnews$subrf > 0 ],
               probs = c(0.33), na.rm = T )

p66 <- quantile( foxnews$subrf[ foxnews$subrf > 0 ],
               probs = c(0.66), na.rm = T )

## Fox News by level of exposure
foxnews$foxnews2000.v2 <- ifelse(foxnews$subrf >= p66, 2,
                              ifelse( foxnews$subrf > 0
                                      & foxnews$subrf <= p33,
                                      1, 0))

## Table of counts
table(foxnews$foxnews2000.v2)

##
##    0    1    2
## 8532 613 650

## Difference in Differences
diff.t <- mean(foxnews$voteshare.diff[foxnews$foxnews2000.v2 == 2],
              na.rm = T)
diff.c <- mean(foxnews$voteshare.diff[foxnews$foxnews2000.v2 == 0],
              na.rm = T)
did.high.nofox <- diff.t - diff.c
did.high.nofox

## [1] -0.01916827
```

```

diff.t2 <- mean(foxnews$voteshare.diff[foxnews$foxnews2000.v2 == 1],
               na.rm = T )
diff.c2 <- mean(foxnews$voteshare.diff[foxnews$foxnews2000.v2 == 0],
               na.rm = T )
did.low.nofox <- diff.t2 - diff.c2
did.low.nofox

```

```
## [1] -0.002484412
```

The diff-in-diff estimate for the ‘Low exposure’ treatment is only -0.00248, which is considerably lower, in absolute terms, than the estimate calculated in the previous question. However, the estimate for the ‘High exposure’ treatment is slightly greater in absolute terms, measuring at -0.01917. Notice that the direction of the effect is consistent across the three estimates. This serves as a further check on the validity of our analysis, confirming the earlier conclusion that exposure to Fox News marginally reduced the Republicans’ share of the vote in the 2000 Presidential election, as compared with the 1996 election.

Question 6

Finally, we consider the effect of having access to Fox News (as measured by `foxnews2000`) on the Republican vote share for each state. Repeat the analysis you have done in Question 4 for each state and compute a state-specific estimate of the Fox News effect (whenever possible). Create a histogram of state-specific effects to examine how much the magnitude of the Fox News effect varies across states. Interpret the results. Finally, compare the average effect across states with the estimate you obtained in Question 4. What does this comparison suggest about the validity of the assumption made in Question 4?

Answer 6

```

## Mean for Towns that received Fox News
fox.st <-
  tapply(foxnews$voteshare.diff[foxnews$foxnews2000 == 1],
        foxnews$state[ foxnews$foxnews2000 == 1], mean,
        na.rm = TRUE)

## Mean for Towns that did not received Fox News
nofox.st <- tapply(foxnews$voteshare.diff[foxnews$foxnews2000 == 0],
                 foxnews$state[foxnews$foxnews2000 == 0],
                 mean, na.rm = TRUE)

diff.by.state <- fox.st - nofox.st[names(nofox.st)!="Nh"]
diff.by.state

```

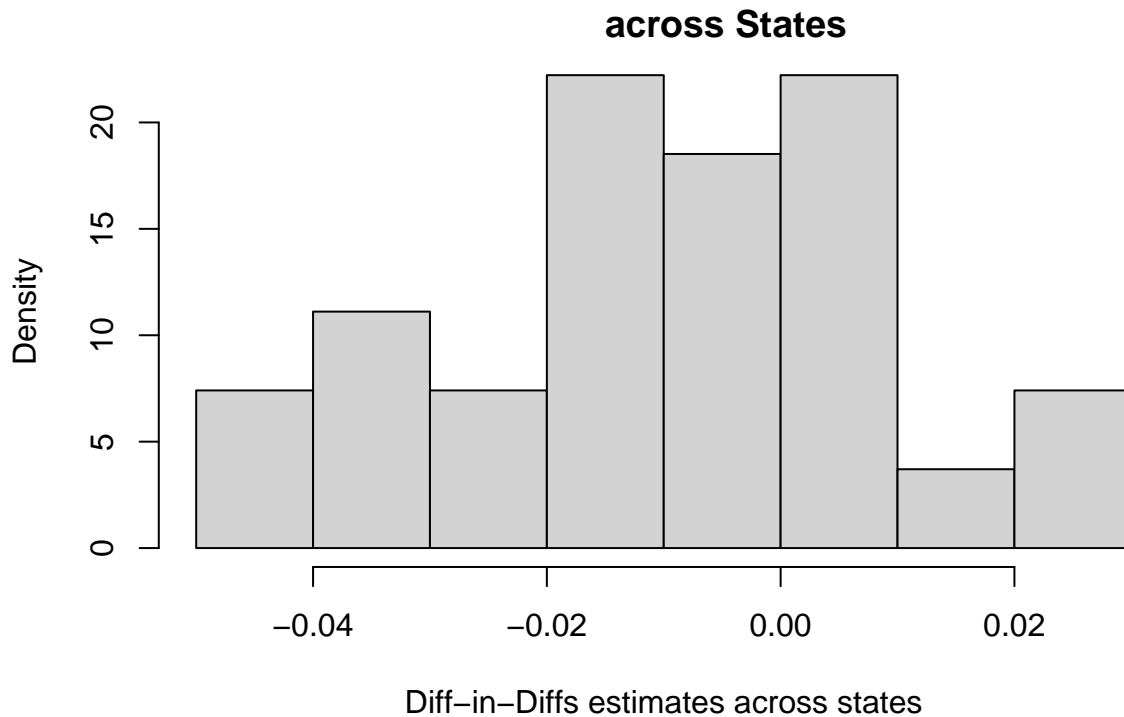
```

##           Ak           Al           Ar           Ca           Ct
## -0.0333012693  0.0296800967 -0.0171267549  0.0032282080 -0.0162569860
##           Hi           Ia           Id           Ma           Me
##  0.0256087543 -0.0093566449 -0.0242790157 -0.0156556278 -0.0458419385
##           Mi           Mn           Mo           Mt           Nd
##  0.0058233618 -0.0227238015 -0.0372781795 -0.0129734597  0.0106292793
##           Nj           Ny           Oh           Pa           Ri
## -0.0004142064  0.0052745892 -0.0181676266 -0.0044929870 -0.0037777697
##           Sc           Tn           Ut           Va           Vt
## -0.0384330797 -0.0096598853  0.0063940299 -0.0491803776  0.0079530270
##           Wi           Wy
## -0.0145332341  0.0001905089

```

```
hist(diff.by.state, freq = FALSE,
     main = "Distribution of the Diff-in-Diffs estimates \n
           across States",
     xlab = "Diff-in-Diffs estimates across states")
```

Distribution of the Diff-in-Diffs estimates



```
comparison <- rbind(DiD,mean(diff.by.state, na.rm = TRUE))
rownames(comparison) <- c("DiD from Question 4",
                          "Average DiD across States")

comparison
```

```
##           [,1]
## DiD from Question 4 -0.01298383
## Average DiD across States -0.01032115
## We can also use a 'for loop' to solve this
## question though tapply() is better
```

```
state.names <- unique(foxnews$state)
diff.by.state.v2 <- rep(NA, length(state.names))
length(state.names)
```

```
## [1] 28
```

```
for(i in 1:length(state.names)) {
  control.group <-
    mean(foxnews$voteshare.diff[foxnews$foxnews2000 == 0 &
                                foxnews$state == state.names[i]],
        na.rm = TRUE)
```

```

treat.group <-
  mean(foxnews$voteshare.diff[foxnews$foxnews2000 == 1 &
      foxnews$state == state.names[i]],
      na.rm = TRUE)
diff.by.state.v2[i] <- treat.group - control.group
}
names(diff.by.state.v2) <- state.names
diff.by.state.v2

##           Ak           Al           Ar           Ca           Ct
## -0.0333012693  0.0296800967 -0.0171267549  0.0032282080 -0.0162569860
##           Hi           Ia           Id           Ma           Me
##  0.0256087543 -0.0093566449 -0.0242790157 -0.0156556278 -0.0458419385
##           Mi           Mn           Mo           Mt           Nd
##  0.0058233618 -0.0227238015 -0.0372781795 -0.0129734597  0.0106292793
##           Nh           Nj           Ny           Oh           Pa
##           NaN -0.0004142064  0.0052745892 -0.0181676266 -0.0044929870
##           Ri           Sc           Tn           Ut           Va
## -0.0037777697 -0.0384330797 -0.0096598853  0.0063940299 -0.0491803776
##           Vt           Wi           Wy
##  0.0079530270 -0.0145332341  0.0001905089

mean(diff.by.state.v2, na.rm = TRUE)

## [1] -0.01032115

```

As it can be noted above, the effect was not homogenous across states. There were some states that experienced an increase in the vote share in favor of the Republicans e.g., Alabama, Hawaii, North Dakota, Vermont; while for some other states like California, New York, and Wyoming, the effect was really small. The rest of the states experienced a decrease in their support for the Republicans. Note that in the case of New Hampshire, there are only towns in the control group, which makes it not possible to estimate the effect of Fox News on the Republicans' vote share. That is why the difference-in-differences estimate from Question 4 is slightly different from the average difference-in-differences estimate across states - in Question 4 that information was used, while in this question it is not possible to use it.