

**Supplementary information to:****Lessons from first campus carbon-pricing scheme (Comment in *Nature* 551, 27–29; 2017)**

doi:10.1038/551027a

*Kenneth Gillingham, Stefano Carattini and Daniel Esty*

The raw data for the analysis can be found at:

[www.environment.yale.edu/gillingham/Raw\\_data\\_for\\_YaleCC.zip](http://www.environment.yale.edu/gillingham/Raw_data_for_YaleCC.zip)

**Data**

This section details the data used to produce Figure 1, which shows the effect of Yale internal carbon pricing scheme on energy consumption. Five variables are used to produce Figure 1:

- A unique identifier for each building (unit) in the Yale campus
- A variable assigning each unit to the control group (value 0), the information treatment (value 1), or the pilot treatment group 2 “carbon price + target” (value 2)
- Monthly information on energy consumption at the unit level, measured in MMBTU
- Dates
- An indicator variable taking the value one during the pilot period

We compare energy consumption during the pilot period with energy consumption during the same period in the previous year. The treatment takes place between December 2015 and May 2016. The pre-period goes from December 2014 to May 2015. The control group is composed of 299 buildings and these buildings received no treatment whatsoever. Five buildings were randomly assigned to the information treatment, which consisted of a monthly energy report explaining what the carbon charge would be, but there were no financial repercussions. The goal of this treatment was to test whether it was simply providing information about the carbon charge that led to a change in behavior (a “salience” effect).

Five units were randomly assigned to the “carbon price + target” pilot treatment group. This group received information provision through a monthly energy report, a carbon price of \$40/ton of carbon dioxide applied to the monthly energy bills based on the carbon intensity of the energy consumed, and an explicit target in the carbon pricing materials of 1% below the baseline emissions. The revenues from the carbon charge were returned to the treated units based on each building’s individual performance relative to the target. For example, consider three buildings: A, B, and C. Suppose building A reduces emissions by exactly 1%. Then at the end of the pilot period the unit will receive back exactly what was paid into the carbon charge fund. In contrast, suppose building B does not change emissions at all from the baseline emissions. Then, the unit will receive back less at the end of the pilot than they were charged throughout the year. Finally, suppose building C reduces emissions more than the 1% target reduction. Then at the end of the pilot period, the unit will receive back more from the carbon charge revenues than they were charged throughout the year. This approach does not assure revenue-neutrality to the university, as there could be a deficit or surplus depending on whether the treated buildings meet or exceed the 1% target on average. Revenue-neutrality would only occur if the buildings collectively meet the 1% reduction target. In the case of the Yale pilot university, units exceeded the target and the university ran a small deficit.

The buildings were randomized into the different treatment groups through a random number generator applied to the pool of all eligible buildings. All of the buildings within each treatment group received the same treatment. 8 out of 10 of the buildings are centrally-supported buildings, meaning that these buildings are not responsible for their own fundraising and utilities expenses. These are the largest fraction of buildings at Yale. Two of these buildings are residential colleges: Berkeley College and Pierson College. Two out of the 10 buildings (one in each treatment group) are self-supported units, meaning that

they are responsible for their own budgeting and fundraising. These buildings are energy ratepayers. Centrally-supported and self-supported units received the same carbon charge and both had to find money in their budgets to pay for this carbon charge. Because the self-supported units are slightly different, we also examine the differences in means excluding these buildings, and find negligible differences in the results.

We further explore the differences in characteristics between the buildings to assure that the randomization was performed properly. We focus on comparing the carbon charge + target group with the control group. Table S1 provides the summary statistics for key variables for the control group and the carbon charge + target treatment group over the pre-period before the treatment. For example, the total energy consumption includes all energy consumption by the building, while the next three energy consumption variables contain three of the largest categories of energy consumption. Table S1 reveals that there are some differences on average between the buildings—as would be expected. However, none of these differences are statistically significantly different from zero to any reasonable level of statistical significance. In fact, all of the p-values of a two-sided t-test are greater than 0.4.

Table S1. Summary statistics for pre-period energy consumption and square footage.							
Target treatment		Total energy consumption (MMBtu)	Energy: steam (MMBtu)	Energy: water (MMBtu)	Energy: electricity (MMBtu)	Emissions (tons CO <sub>2</sub> )	Square footage (ft <sup>2</sup> )
0	Mean	2358.42	373.74	293.18	104.87	151.30	49100.94
(Control group)	SD	13441.07	848.92	785.38	292.01	829.58	73348.23
	Min	0.053	0	0	0	0.005	0
	Max	157823.2	6161.67	5551.85	2719.03	9832.455	535971
	N	299	299	299	299	299	299
1	Mean	1059.85	506.61	205.23	74.53	78.96494	75438.4
(Target treatment)	SD	1471.35	743.466	356.22	138.08	106.7383	69787.28
	Min	38.33412	0	0	0	2.810553	7545
	Max	3471.455	1644.713	822.65	320.18	253.2174	173012
	N	5	5	5	5	5	5
		P-value for t-test					
H <sub>a</sub> : diff != 0		0.829	0.728	0.803	0.817	0.846	0.426
Notes: All variables are based on the pre-period from December 2014 to May 2015. The energy consumption and emissions variables are per month. SD refers to the standard deviation, N refers to the number of buildings. The p-value is given for a two-sided t-test of differences in means, with a null hypothesis that the difference is zero.							

The data used in this study are made available by the Yale Carbon Charge Project (<http://carbon.yale.edu/>), the organizational structure within Yale that manages the internal carbon pricing scheme. The Yale Carbon Charge Project received data on energy consumption from the Yale University Office of Facilities, which manages building-level metering for the university campus.

## Analysis

The analysis focuses on the “carbon price + target” pilot arm (treatment group 2) and the information treatment (treatment group 1), which are compared to a control group. Treatment groups 3 and 4 are not included in this illustrative analysis for simplicity and conciseness in the message. There is little added insight from including the additional treatment groups for illustrating how an internal carbon price works in practice.

The analysis behind Figure 1 is composed of two stages:

Stage 1: For each unit in the sample the difference between the average energy consumption during the pilot period and the average energy consumption during the pre-period is computed; this difference is divided by the pre-period level to obtain the percentage change in energy consumption.

Stage 2: Unit-level differences are averaged out at the treatment level and standard errors are computed to provide a sense of the dispersion. In the spirit of a difference-in-differences approach, we can isolate the variation due to treatments 1 and 2 by netting out the variation that is also observed in the control group, which we attribute to a warmer winter than in previous year. Hence, Figure 1 shows the average effect of the pilot treatments on energy consumption, net of the variation observed in the control group. In Figure 1, all changes are reported in comparison to the control group.

Figure 1 takes the results of this simple analysis to compare the average difference in energy consumption between the pre-period and the pilot period for the “carbon price + target” pilot arm, the information treatment, and the control group.

Figure 1 is produced in Excel. The average difference in energy consumption and the standard errors for each treatment are computed using the statistical software Stata.

Figure S.1 below shows the average effect of the pilot treatments on energy consumption, before netting out the variation observed in the control group.

Finally, a t-test is computed to test the statistical significance of the differences in means between each treatment group and the control group. The F-statistic indicated that a t-test allowing for unequal variances was appropriate for the comparison of means between the target + pilot group and the control group. We thus perform a t-test allowing for unequal variances (sometimes called “Welsh’s t-test”). We examine both a one-sided and two-sided t-test, and report the results from the one-sided t-test because we are more interested in understanding whether the energy savings from the treatment group are greater than the energy savings in the control group.

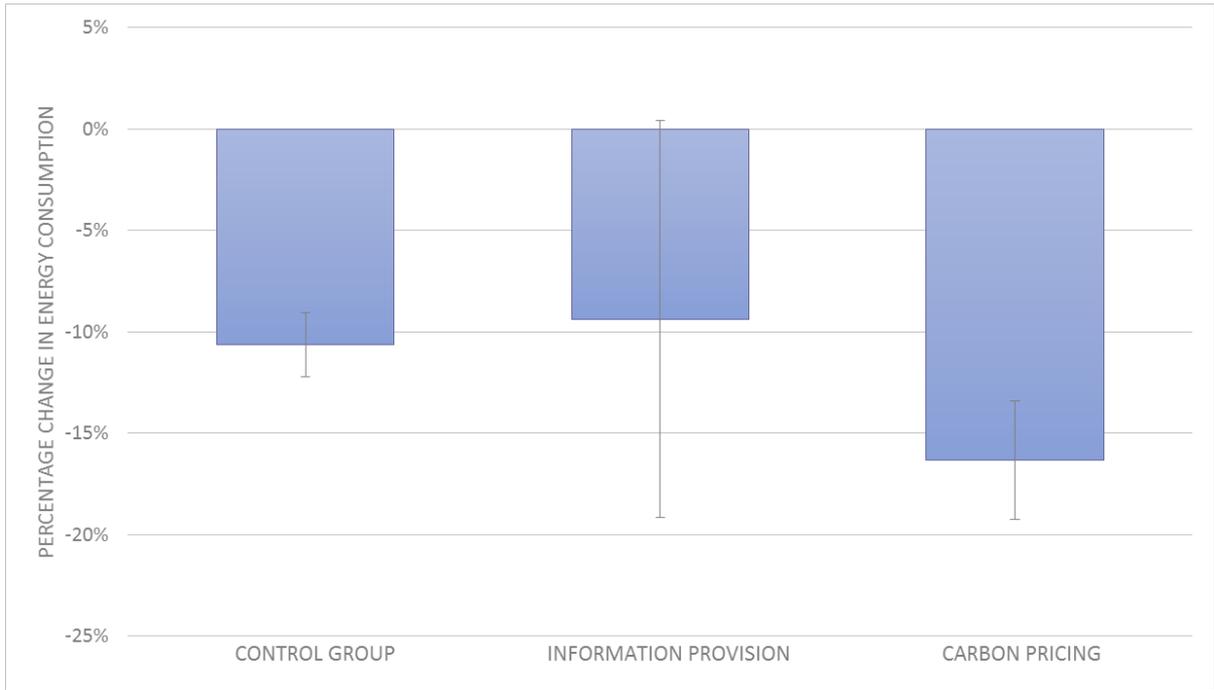


Figure S.1. The effect of the Yale internal carbon pricing December 2015-May 2016 pilot effort. For both the treatment and control, we compare the energy consumption during the pilot to a comparable pre-period December 2014-May 2015. The whisker bars indicate  $\pm$  one standard error.

### Emissions as outcome variable

Figure S.1 shows the average effect of the pilot treatments on energy consumption. Figure S.1 can be reproduced using CO<sub>2</sub> emissions as outcome variable. Values for CO<sub>2</sub> emissions are provided directly by the Yale Carbon Charge Project based on energy sources and Yale-specific emissions factors. Table S.1 provides descriptive statistics for this variable for the control group and target treatment.

Figure S.2 shows the average effect of the pilot treatments on CO<sub>2</sub> emissions (before netting out the change in emissions from the control group).

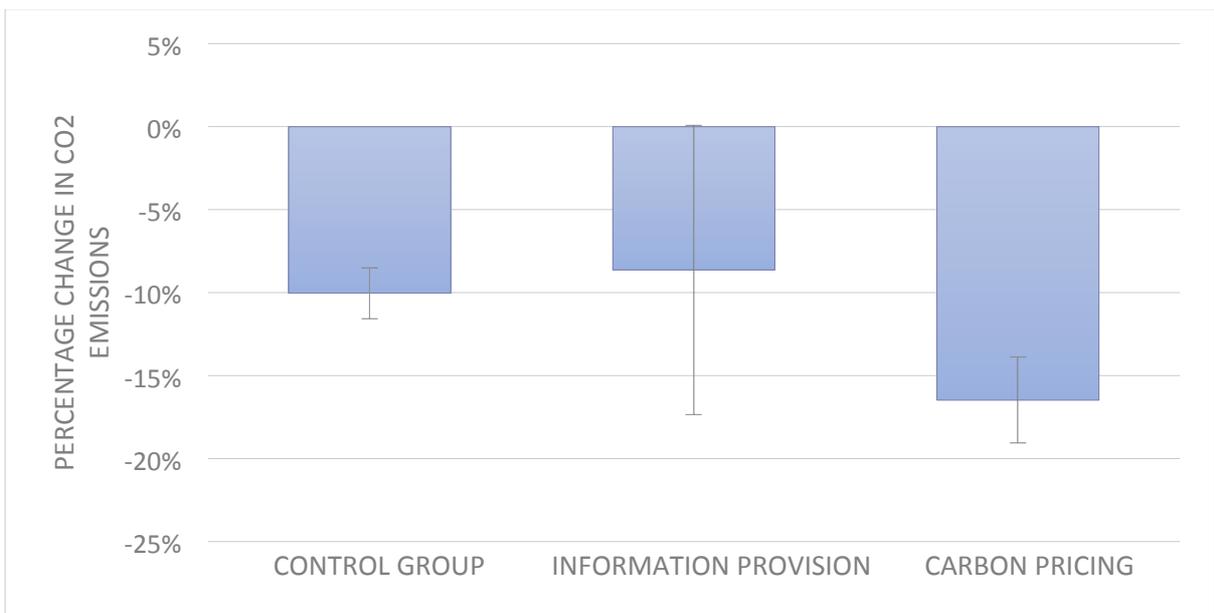


Figure S.2. The effect of the Yale internal carbon pricing December 2015-May 2016 pilot initiative. For both the treatment and control, we compare CO<sub>2</sub> emissions during the pilot initiative to a comparable pre-period, December 2014-May 2015. The whisker bars indicate  $\pm$  one standard error.

Using CO<sub>2</sub> emissions as outcome variable does not change our conclusions about the effect of the pilot treatments. As with Figure S.1, all three groups had lower emissions due to the warmer winter, and the decline is still 5 percentage points larger for the internal carbon pricing group than the control group (one-sided t-test for the carbon pricing decline being larger p-value= 0.034).

### Stata do file code

Energy consumption as outcome variable:

```
clear all

use Raw_data.dta

*Stage 1

bysort pilot_period id: egen pilotenergy=mean(energy)

keep if time==660 | time==675

xtset id pilot_period

gen d_pilotenergy=d.pilotenergy
gen d_relative_pilotenergy=d_pilotenergy/l.pilotenergy

#delimit ;
global diff "
d_relative_pilotenergy
";
#delimit cr*/

*Stage 2 (0 "Control group"; 1 "Information"; 2 "Target")

foreach var in $diff {
preserve
collapse (mean) `var' (semean) `var'_sem=`var', by(treatment)
export excel using Stage2.xls, replace
restore
}

*T-test

gen information=(treatment==1)
replace information=. if treatment==2
gen target=(treatment==2)
replace target=. if treatment==1
```

```
sdtest d_relative_pilotenergy, by(information)
ttest d_relative_pilotenergy, by(information)
```

*\* The F-stat was used to determine whether to assume equal or unequal variances for the t-test.*

```
sdtest d_relative_pilotenergy, by(target)
ttest d_relative_pilotenergy, by(target) unequal
```

*\*Refer to the Excel file for the generation of Figure 1*

*\* Code for the balance of covariates analysis included in this SM:*

```
clear all
```

```
use Raw_data_covariates.dta
```

```
gen information=(treatment==1)
replace information=. if treatment==2
```

```
gen target=(treatment==2)
replace target=. if treatment==1
```

```
tabstat energy sumst sumw sumel sumem sq, statistics(mean sd min max n) by(information)
tabstat energy sumst sumw sumel sumem sq, statistics(mean sd min max n) by(target)
```

```
#delimit ;
global covariates "
energy sumst sumw sumel sumem sq
";
#delimit cr*/
```

```
foreach var in $covariates {
ttest `var', by(information)
}
```

```
foreach var in $covariates {
ttest `var', by(target)
}
```

CO<sub>2</sub> emissions as outcome variable:

```
clear all
```

```
use Raw_data_emissions.dta
```

*\*Stage 1*

```
bysort pilot_period id: egen pilotemissions=mean(emissions)
```

```
keep if time==660 | time==675
```

```
xtset id pilot_period
```

```
gen d_pilotemissions=d.pilotemissions
```

```
gen d_relative_pilotemissions=d_pilotemissions/l.pilotemissions
```

```
#delimit ;
```

```
global diff "
```

```
d_relative_pilotemissions
```

```
";
```

```
#delimit cr*/
```

```
*Stage 2 (0 "Control group"; 1 "Information"; 2 "Target")
```

```
foreach var in $diff {
```

```
preserve
```

```
collapse (mean) `var' (semean) `var'_sem=`var', by(treatment)
```

```
export excel using Stage2_emissions.xls, replace
```

```
restore
```

```
}
```

```
*T-test
```

```
gen information=(treatment==1)
```

```
replace information=. if treatment==2
```

```
gen target=(treatment==2)
```

```
replace target=. if treatment==1
```

```
sdtest d_relative_pilotemissions, by(information)
```

```
ttest d_relative_pilotemissions, by(information)
```

```
* The F-stat was used to determine whether to assume equal or unequal variances for the t-test.
```

```
sdtest d_relative_pilotemissions, by(target)
```

```
ttest d_relative_pilotemissions, by(target) unequal
```

```
*Refer to the Excel file for the generation of Figure 1
```

## **Further Details on Key Findings from the Yale Carbon Pricing Project**

The article refers to an internal report providing a set of illustrative findings from the Yale carbon pricing experiment, which complement our analysis of the pilot's effectiveness. This report is publicly accessible at: [http://carbon.yale.edu/sites/default/files/files/Carbon\\_Charge\\_Pilot\\_Report\\_20161010.pdf](http://carbon.yale.edu/sites/default/files/files/Carbon_Charge_Pilot_Report_20161010.pdf).

We report here some relevant information underlying these additional findings.

1. The pilot increased understanding of energy use among treated units.

The pilot was accompanied by semi-structured interviews, an exit survey and treatment-specific focus groups. The aim was to examine the levels of understanding, motivation, and action among units involved in the pilot (no surveys were conducted with units in the control group). In the exit survey, building appointees to the pilot were asked a variety of questions about their current level of understanding of energy use and how and why it changed from the pre-treatment level. Changes were scored “higher” (1), “the same” (2), and “lower” (3). 75% of the “target” pilot arm report a higher understanding of and attention to energy use data. Overall, more than 2/3 of participants report higher understanding and attention.

### **ACKNOWLEDGEMENTS**

The authors would especially like to thank Ryan Laemel and Casey Pickett for providing access to the Yale Carbon Charge Project data and valuable background information on the pilot charge. We would also like to thank William Nordhaus and Sharon Oster for constructive suggestions. Carattini acknowledges funding from the Swiss National Science Foundation, grant P2SKP1\_165028, and from the Centre for Climate Change Economics and Policy, which is funded by the UK Economic and Social Research Council.