Unclassified **ENV/WKP(2012)10**

Organisation de Coopération et de Développement Économiques
Organisation for Economic Co-operation and Development **05-Dec-2012**

English - Or. English

ENVIRONMENT DIRECTORATE

ENV/WKP(2012)10
Unclassified

Environment Working Paper No. 51

TESTING THE EFFECT OF DEFAULTS ON THE THERMOSTAT SETTINGS OF OECD EMPLOYEES

by Zack Brown*, Nick Johnstone, Ivan Haščič, Laura Vong, and Francis Barascud

** Corresponding author*
Email: Zack.BROWN@oecd.org
Telephone: +(33-1) 45 24 18 32

Keywords: Behavioural economics, energy efficiency, field experiments

JEL codes: B5, C1, C9, H3, Q4

All Environment Directorate Working Papers are available at
www.oecd.org/environment/workingpapers

JT03332309

Complete document available on OLIS in its original format

English - Or. English

ABSTRACT

Default options have been shown to affect behaviour in a variety of economic choice tasks, including health care and retirement savings. Less research has tested whether defaults affect behaviour in the domain of energy efficiency. This study uses data from a randomized controlled experiment in which the default settings on office thermostats in an OECD office building were manipulated during the winter heating season, and employees' chosen thermostat setting observed over a 6 week period. Using difference-in-differences, panel, and censored regression models (to control for maximum allowable thermostat settings), we find that a 1°C decrease in the default caused a reduction in the chosen setting by 0.38°C on average. Sixty-five percent of this effect could be attributed to office occupant behaviour (p-value=0.044). The difference-in-differences model shows that small decreases in the default (1°) led to a greater reduction in chosen settings than large decreases (2°). We also find that office occupants who are more apt to adjust their thermostats prior to the intervention were less susceptible to the default. We find no evidence that offices with multiple occupants displayed different patterns in thermostat choices than single-occupant offices. We conclude that this kind of intervention can increase building-level energy efficiency, and discuss potential explanations and broader policy implications of our findings.

Keywords: Behavioural economics, energy efficiency, field experiments

JEL codes: B5, C1, C9, H3, Q4

RESUME

Il a été démontré que les options par défaut influaient sur le comportement dans diverses situations de choix économique, portant par exemple sur le système de santé ou le régime de retraite. Cependant, l'incidence des options par défaut sur le comportement dans le domaine de l'efficacité énergétique a fait l'objet de travaux de recherche moins nombreux. Pour cette étude, des données ont été recueillies dans le cadre d'une expérience aléatoire contrôlée ayant consisté à manipuler le réglage par défaut des thermostats installés dans les bureaux d'un bâtiment de l'OCDE pendant la période de chauffage hivernale, et à observer le réglage choisi par les salariés sur une période de 6 semaines. Des modèles fondés sur la méthode des « différences de différences », des données de panel et une analyse de régression censurée (prenant en compte les réglages thermostatiques maximum admissibles) permettent de constater qu'une baisse de la température par défaut de 1°C se traduit par une réduction de 0.38°C en moyenne de la température choisie. Soixante-cinq pour cent de cet effet pourrait être attribué au comportement de l'occupant du bureau (valeur-p=0.044). Le modèle de « différences de différences » montre qu'une légère baisse de la température par défaut (1°) entraîne une plus forte réduction de la température choisie qu'une baisse importante (2°). Nous constatons aussi que les occupants des bureaux les plus enclins à ajuster leur thermostat avant l'intervention ont été moins sensibles au réglage par défaut. Nous ne trouvons pas de différence quant aux choix de température entre les bureaux occupés par plusieurs personnes et les bureaux individuels. Nous concluons que ce type d'intervention peut accroître l'efficacité énergétique au niveau des bâtiments, et examinons les explications possibles et les enseignements plus généraux qui peuvent être tirés de nos résultats pour l'élaboration des politiques publiques.

Mots clés: Économie comportementale, efficacité énergétique, expériences de terrain

Classifications JEL: B5, C1, C9, H3, Q4

FOREWORD

This report has been prepared by Zack Brown, Nick Johnstone, Ivan Haščič, Laura Vong, and Francis Barascud (OECD). It is a contribution to the OECD Environment Directorate project on "Managing the Transition in Environmental Policy Reform" (www.oecd.org/environment/behaviour.htm). The support of Peter Lübkert (Head of the Buildings, Logistics and Services Division, OECD) is gratefully acknowledged.

A version of this report was presented at the June 2012 meeting of the OECD Working Party on Integrating Environmental and Economic Policies and it has benefited from the comments received. It represents the views of the authors, and not necessarily those of the OECD or its member countries. This paper is released as part of the OECD Environment Working Paper series. It can be downloaded from the OECD website (www.oecd.org/env/workingpapers).

TABLE OF CONTENTS

EXECUTIVE SUMMARY

Default options have been shown to affect behaviour in a variety of economic choice tasks, including health care and retirement savings. By default option, we mean an alternative that is automatically selected when an agent makes no active choice. Although standard economic theory implies that such defaults should not determine choices when there is no cost associated with switching from one option to another, in practice numerous scientific demonstrations have been conducted in which default settings do evidently influence agents' choices: Well-known previous examples have demonstrated that employee selection of retirement plans can be strongly influenced by the default option, and – in the context of organ donation – that an opt-out default for organ donor programs can double participation rates relative to an opt-in default or no default.

Less research has tested whether defaults affect behaviour in the domain of environmental policy. This is in spite of the fact that there are many examples of default options in choices relating to the environment. These include opt-in/opt-out defaults for electronic/paperless bank statements, default options for carbon offsets associated with air travel, as well as an array of additional examples in the context of consumer energy and water use.

This study uses data from a randomized controlled experiment in which the default settings on office thermostats in an OECD office building were manipulated during the winter heating season, and employees' chosen thermostat setting observed over a 6 week period. Occupants in the 1st floor of the OECD's Marshall Building were randomly assigned to one of three treatments: (I) a control group in the thermostat default was left at its original setting of 20°C, (II) a group in which the default was uniformly decreased 1°C per week until a default setting of 17°C was reached, and (III) a group in which the default was increased to 21°C in the 1st week before being decreased by 1°C per week until reaching a setting of 18°C. The changes in default settings were not announced to the occupants. Occupants' chosen thermostat settings, along with office occupancy and measured temperature, were then monitored remotely via the computerized building heating and cooling control system.

Results indicate that small decreases in the default (1°) led to a greater reduction in chosen settings than large decreases (2°). We also find that office occupants who are more apt to adjust their thermostats prior to the intervention were less influenced by changes in the default. The direct policy message from this small experiment is limited in scope, but clear: Small reductions in the defaults of office thermostats can lead to lower temperature settings by occupants in the winter heating season, which when scaled up to the whole building should translate into lower energy use. However, if the reduction in default temperature is too large, then occupants respond actively, increase their temperature settings, over-riding the effects of the change in the default setting.

More broadly, this study raises questions about why defaults appear to "work." In this case, the intervention used here most likely took advantage of occupants' insensitivity to small changes in temperature, or their aversion to expending cognitive effort worrying about the temperature in their offices. If these explanations are right, then the welfare implications of these types of interventions deserve further analysis. As behavioural interventions become more popular in public policy toolboxes, the individual welfare implications of these interventions will become a more frequent topic of discussion. Consequently, it will be important to understand in more detail the mechanisms underlying individuals' responses to these interventions.

What are the more specific policy implications of the results? It is important to emphasise that such measures cannot be considered as substitutes for more direct measures such as price-based instruments and direct forms of regulation. However, they can be a useful complement. In the case under analysis this is likely to be particularly true in the presence of split incentives or agency control problems which characterise a large number of relevant agents (tenants, office workers, etc.). On the one hand, the incentives of a price-based measure will not be 'transmitted' to those undertaking decisions in terms of actual consumption. On the other hand, sole reliance on a direct form of regulation may constrain their capacity to make decisions which reflect their actual temperature preferences. The use of a behavioural 'prompt' alongside other measures - in this case a default setting – can result in reduced energy consumption in a manner which is cost-effective and consistent with utility maximisation.

TESTING THE EFFECT OF DEFAULTS ON THE THERMOSTAT SETTINGS OF OECD EMPLOYEES

Introduction

There are many documented cases in which the default option in a choice task can strongly influence the behaviour of economic agents. By default option, we mean an alternative that is automatically selected when an agent makes no active choice. Although standard economic theory implies that such defaults should not determine choices when there is no cost associated with switching from one option to another, in practice numerous scientific demonstrations have been conducted in which default settings do evidently influence agents' choices: The most famous examples include the demonstration that employee selection of retirement plan is strongly determined by the default option (Madrian & Shea, 2001), and an experiment showing that an opt-out default for organ donor programs can double participation rates relative to an opt-in default or no default (Johnson & Goldstein, 2003). Of course, a number of cultural factors, sociological norms, and individual experience can mediate the impact of default options, posing important questions about the external validity of such experiments. But there is now little question that defaults matter, and that this can have implications for public policy.

However there are still few examples of scientific studies which examine the role of default options in environmental policy. While there are numerous instances of default options for a variety of environmentally related choice tasks—opt-in/opt-out defaults for electronic/paperless bank statements, default options for carbon offsets associated with air travel, as well as an array of additional examples in the context of consumer energy and water use—only a few of these are covered by peer-reviewed studies: Pichert and Katsikopoulos (2008) find evidence from a series of observational studies and laboratory experiments that switching the default selection of electricity source to a more "green" (i.e. less polluting) option can significantly affect the adoption rates of green electricity among consumers. Löfgren et al. (2012), in a study which offered CO_2 offsets for airline travel to a sample of economists, find that "experienced" subjects are less likely to be influenced by default options. This latter study suggests that defaults can play a particularly important role in the protection of environmental goods and services, a domain in which individuals typically have less experience as compared to commonly exchanged goods.

A variety reasons have been proposed to explain why defaults appear to affect choices. The endowment effect—that is, the notion that people feel that they lose more from giving up a good than they gain by being granted the same good—is sometimes used to explain the effects of defaults (*e.g.*, Marzilli et al., 2011). Yet, many examples of defaults, such as paperless bank statements and decisions about thermal comfort in buildings, involve nothing which we may construe as representing a real endowment. In these situations, other hypotheses are necessary to explain why defaults appear to "work." Such explanations include the notion that defaults influence decisions by drawing more attention to the default than to other options (e.g. Johnson et al., 2002), that people perceive recommendations implicit in defaults (McKenzie et al., 2006), and observations that actively switching a choice from the default, as opposed to passively accepting it, can invoke "effortful" reasoning (Kahneman, 2003). People may choose not to exert this effort, obviating a reasoned choice between all available alternatives.

Understanding the reasons for why defaults affect environmentally-relevant behaviours in some contexts and not others is necessary if default-setting is to be used for public policy. On the one hand, if the effect of defaults in a given context results from drawing people's attention to a previously unconsidered option, then a change in the default can be a means of overcoming an information failure. On the other hand, if individuals perceive defaults as conveying an implicit policy recommendation (e.g. organ donation) then a change in the default to encourage such behaviour can be considered as a form of moral suasion. And finally, the persistence of the effect of defaults over time—and therefore their relevance for broader policy objectives (e.g. achieving lasting gains in energy efficiency)—is likely to depend on the underlying reason for these effects.

In this paper we assess the effects of changes in default settings in a novel application (office thermostat settings). While some might view measures which bring about behavioural change through changes in default settings as paternalistic (e.g. Mitchell, 2005), such a view is not warranted in our case as long as the adjusted setting is 'reasonable'. After all, the default has to be set somewhere.

Even if changes in default settings are considered to be legitimate policy interventions, their effectiveness needs to be evaluated. For example, it is by no means evident that they are a more efficient means to overcome information failures than explicitly drawing attention to different options. Similarly, if they are to be used implicitly as a means of moral suasion, policymakers may do better by making this recommendation explicit. Although the research described above has shown defaults to be effective in discrete choice settings (*e.g.* either opt to be an organ donor or not), much less work has looked at the role of defaults in settings of continuous choice, which is essentially the case when selecting a thermostat temperature.

To analyse the role of defaults in an environmentally relevant field setting, we conducted a simple experiment among employees at the OECD to see how the default thermostat setting affected occupants' chosen temperature settings over a 6 week period. The basic research question for the experiment was: How much does changing the default setting on office thermostats affect the chosen thermostat settings in offices? Each office in the experiment included a thermostat that was accessible to the occupants and for which the default setting was 20°C prior to the experiment. This meant that, unless occupants actively changed their thermostat settings, the heating, ventilation, and air condition (HVAC) system would attempt to drive the office temperature to 20°C.

Analysis of data from the experiment indicates that building managers would achieve lower energy use by decreasing the winter default settings on office thermostats from 20°C to 19°C. However, the data indicate that decreasing the default setting by more than 2° would cause occupants to actively intervene and increase their temperature setting to what prevailed with the 20° default, thereby yielding no decrease in energy use. Possible reasons for why occupants did not actively respond to smaller changes in defaults are discussed at the end of this article.

Materials and Methods

A randomized controlled experimental design was used in the study as follows: All 93 occupied offices on the first floor (above ground floor) of the OECD's Marshall Building were selected for inclusion in the experiment. Consisting of 29,295 square meters with 6 floors and 2 basement levels, the Marshall Building was constructed in 1969, renovated in 2008, and now accommodates 1,105 employees throughout its approximately 608 offices. The heating cooling and air conditioning (HVAC) system is controlled via a computerized system installed by Honeywell®. This system permits building engineers to remotely control a number of HVAC parameters for specific offices, and to monitor both the temperature and user-specified thermostat settings in each of the offices. The parameter for the default thermostat setting was manipulated experimentally. Under baseline conditions, the default is set at 20°C during the winter heating season and occupants are permitted to adjust the thermostat upwards to a maximum of 3°C above the default.

To collect baseline data, thermostat settings, measured temperature, and office occupancy (whether or not someone was present) were recorded for a 1 week period, during OECD's working hours (Table 1). Thermostat settings and temperature were recorded hourly, and occupancy was recorded twice daily, once in the morning and again in the afternoon. While baseline data were being collected, two distinct, three-week schedules were formulated for changes to thermostat defaults. These two schedules (shown in Figure 1) comprised the two "treatment" arms of the experiment. The first treatment simply lowered the default setting by 1°C per week over a three week period, ending at 17°C. The second treatment first increased the default by 1°C (to 21°C) in the first week of the treatment, before lowering the setting by 1°C per week for two subsequent weeks, ending at 19°C. A control group was also specified, in which the default setting remained at 20°C throughout the experiment.

Table 1: Baseline data from a week before the interventions.

N = 82 offices	*Mean*	*Standard Deviation*[2]	*P-value*[3]
Temperature setting (°C)[1]	21.10	1.81	0.72
Temperature sensor (°C)[1]	21.46	1.99	0.60
Occupants per office	1.94	0.95	0.46
Fraction of work hours occupied	0.85	0.28	0.95

[1] *For occupied hours only.*
[2] *Standard errors reported here are for between-office variation in each office's mean value over the baseline period.*
[3] *P-values from a Wald test for regressions of each baseline characteristic on treatment assignment. Insignificant p-values (greater than 0.10) indicate that random assignment of offices to treatments was successful.*

Offices were then randomly assigned to one of the two treatment arms or to the control. To test that the randomization was successful, we computed Wald test statistics from the linear regression $X_i = \beta_0 + \beta_1 q_{1i} + \beta_1 q_{2i} + \epsilon_i$, where X_i is a baseline characteristic and q_{1i}, q_{2i} are the treatment dummy variables. The p-values of these χ^2-statistics are shown in Table 1; that they are all insignificant indicates that the randomization was successful. For each treatment, the default schedules were then implemented by building management, and data on thermostat settings, temperature, and occupancy were collected throughout the three treatment weeks and for two subsequent weeks.

Figure 1: Scheduled changes in defaults for treatment and control groups.

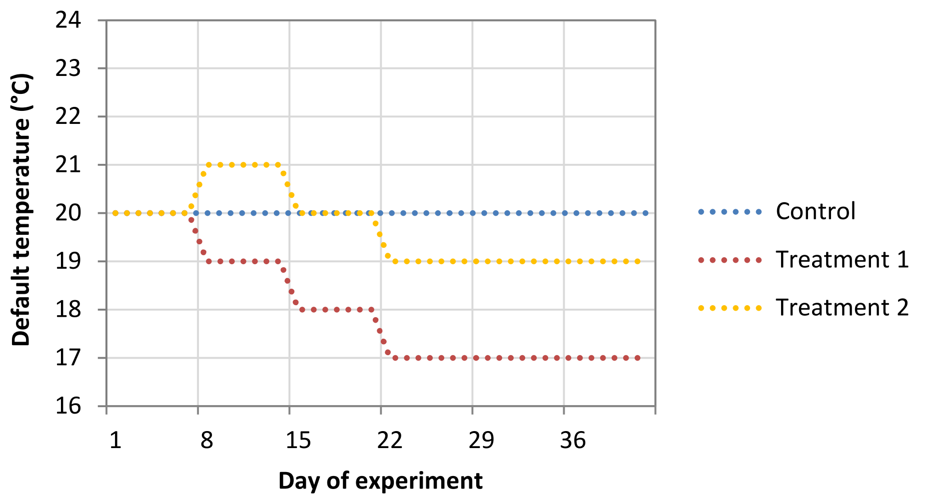

Note: both treatments and the control are the same during the baseline period (days 1-8).

Results

Overview of main findings

The primary outcomes of interest are the temperature settings chosen by office occupants. Figure 2 plots these data over time for each of the treatment and control groups and among only those offices which were actually occupied. The scheduled changes in default settings are also shown for reference. There are two important points contained in this figure, one relating to the validity of the experiment and the last comprising our main result: First, those in the control group increased their thermostat setting greatly during February; this was likely a response to the unusually cold weather during that period (see Figure 3) and certainly was not related to the default setting, since that remained at a constant 20°C throughout.

Most importantly, Figure 2 shows that both Treatment 1 and Treatment 2 average settings paralleled the scheduled changes in default settings (Figure 1) for the majority of the experiment, but not for the last two weeks. This suggests that most of the default changes were not negated by individual responses. Only in the last stage of Treatment 1, where the default setting was decreased to 17°, did individuals seem to finally respond to the external changes by adjusting their settings higher. A notable detail here is that the Treatment 2 group—which ended the experiment with a default 2° higher than Treatment 1—exhibited an average setting that was actually 0.5°C *lower* than Treatment 1, although the difference is not statistically significant at the 5% level.

Figure 2: Changes in thermostat settings over time for treatment and control groups.

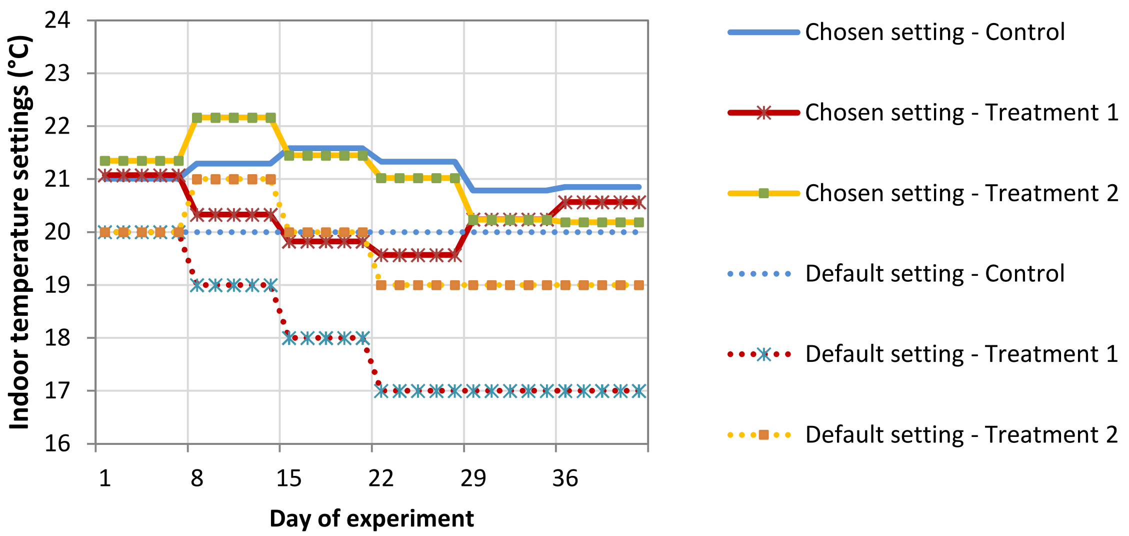

Note that there is overlap in defaults and chosen settings in days 1-8 of the experiment because this was the baseline period of the experiment.

Statistical and econometric analysis

To test the statistical significance of our results, we calculate the average treatment effect (ATE) for each treatment in each week of the experiment. Here, we also address the potential confounder of the maximum office thermostat setting, which is collinear with the experimentally controlled default setting. The maximum thermostat setting could bias our estimates of the behavioural response to the default changes if we do not account for this collinearity.

We use a difference-in-differences (DID) approach to calculate treatment effects for a given treatment in a given week: The ATE for a given week and treatment is the difference in mean temperature settings between the treatment and control group for that week, subtracting the same difference for the baseline week (before treatment was applied). The estimated ATE_{wt} for week w and treatment $t = 1{,}2$ is:

$$ATE_{wt} \equiv (\bar{s}_{wt} - \bar{s}_{wC}) - (\bar{s}_{0t} - \bar{s}_{0C}) \quad (1)$$

where $\bar{s}_{wt}$ is the average temperature setting for week w and treatment t (with $t = C$ denoting the control). The DID approach controls for pre-existing differences between the treatment groups. Even though these differences are not statistically significant due to successful randomization of treatment assignment (Table 1), we use DID to ensure the robustness of our results and for interpretability. The ATEs are estimated via the following linear regression:

$$s_{ih} = \beta_0 + \sum_{w=2}^{6} \beta_w d_{wh} + \sum_{t=1}^{2} \gamma_t q_{ti} + \sum_{w=2}^{6} \sum_{t=1}^{2} ATE_{wt}(d_{wh} \times q_{ti}) + \epsilon_{ih} \quad (2)$$

where s_{ih} is the temperature setting for office i in hour h of the experiment, the d_{wh} are dummy variables for each week w of the experiment, and the q_{ti} are dummy variables for treatment assignment t. The variable ϵ_{ih} is econometric error. We use random and fixed effects models to estimate the parameters in the above regression, clustering the standard errors of the parameter estimates at the office level (Wooldridge, 2010).

Table 2: Average treatment effect on chosen temperature settings, by treatment arm and week.

N = 26,712 office-hours	*Week of experiment*					
Groups = 87 offices	*1*	*2*	*3*	*4*	*5*	*6*
Control (30 offices)						
Default setting (°C)	20	20	20	20	20	20
Percent obs. at upper constraint	21%	28%	39%	32%	21%	20%
Treatment 1 (27 offices)						
Default setting (°C)	20	19	18	17	17	17
Percent obs. at upper constraint	26%	25%	39%	61%	76%	81%
ATE (°C)	-	-1.125***	-1.883***	-1.966***	-0.515	-0.408
	-	(0.371)	(0.454)	(0.464)	(0.451)	(0.498)
Adjusted ATE (°C)	-	-0.884**	-1.200**	-0.313	1.150**	1.194**
	-	(0.367)	(0.466)	(0.539)	(0.502)	(0.510)
Treatment 2 (30 offices)						
Default setting (°C)	20	21	20	19	19	19
Percent obs. at upper constraint	33%	27%	42%	45%	6%	10%
ATE (°C)	-	0.381	-0.569*	-0.770**	-1.001***	-1.030**
	-	(0.294)	(0.338)	(0.318)	(0.343)	(0.421)
Adjusted ATE (°C)	-	0.0914	-0.606*	-0.495	-1.015***	-1.024**
	-	(0.293)	(0.333)	(0.314)	(0.343)	(0.422)

*See text for formal definition of the average treatment effect (ATE), and the Adjusted ATE. The asterisks denote statistical significance of the treatment effect, with ***,**, and * denoting 1%, 5%, and 10% statistical significance, respectively. Robust standard error estimates, in parentheses, are from a random effects panel regression model, calculated via clustering by each of the 87 offices for which a full set of measurements were obtained (and which were occupied during the experiment). The default settings in the shaded rows were experimentally controlled.*

As can be seen from the table, decreasing the default setting in Treatment 1 caused a statistically significant reduction of between 1 and 2 degrees in the chosen thermostat setting for weeks 2 through 4 of the experiment, but the effect dissipated in weeks 5 and 6. Meanwhile, the chosen settings of those exposed to Treatment 2 became significantly lower than those in the control group towards the end of the experiment, from Week 4 onward. This is consistent with the summary results shown in Figure 1.

Recall from the description of the experiment above that the HVAC control system constrains the maximum temperature setting to +3°C above the default setting. This maximum setting automatically adjusted whenever the default changed in the experiment, so that a default setting of 19°C implied a maximum setting of 22°C and so on. Thus, it is possible that lowering the default revealed statistically significant ATEs simply because the upper constraint on settings pulled the average setting down. Indeed, as can be seen in Table 2, the majority of observations are constrained by the upper limit on their temperature setting. To account for this potential downward bias, we calculate an "adjusted" temperature setting variable as follows:

$$\tilde{s}_{ih} \equiv \begin{cases} s_{ih} \text{ if } s_{ih} < (\text{default}_{ih} + 3) \\ 23 \text{ if } s_{ih} \geq (\text{default}_{ih} + 3) \end{cases} \tag{3}$$

In words, this new variable $\tilde{s}_{ih}$ is the same as the observed temperature setting, except when s_{ih} is at its constrained value: When an office's temperature setting is at its maximum value, we define the adjusted temperature setting as the maximum value of 23°C that was available during baseline conditions. We then use this adjusted temperature setting variable as the dependent variable in a regression equivalent to (2). We call the resulting treatment effects "Adjusted ATEs," which are reported in Table 2.

Note that this estimation approach is a conservative way to minimize bias in estimating the effect of decreasing thermostat defaults on behaviour, controlling for the corresponding decrease in the constraint: We do not know exactly what occupants would have selected as their setting if they had not been constrained by their maximum setting. By simply assuming that they would have selected the maximum baseline setting of 23°C and including these adjusted measurements in our sample decreases the likelihood that we would find a significant downward effect of decreasing defaults on chosen settings. In a later section, we take a more sophisticated approach to address this data censoring, using tobit regression models.

Examining the Adjusted ATEs in Table 2, we see indeed that the estimated effects are smaller and less significant than the unadjusted ATE estimates. Nevertheless, the general pattern of effects for the adjusted temperature setting (and statistical significance) remains: As the default decreases by 1 or 2 degrees, the chosen settings decline as well—now controlling for the decreasing maximum setting. However, a drastic decrease in the default by 3°C or more stimulates a significant, active response from occupants.

The role of weather, preference/perception heterogeneity, and multi-person offices

While we have shown above that the treatments had a significant impact on average temperature settings chosen by occupants, here we focus on the potential environmental factors and behavioral heterogeneity which mediate these effects. The environmental factor we focus on is weather—measured by daily average outdoor temperature—since there was a major cold snap that occurred during the middle of the experiment and that could have affected occupants' indoor temperature preferences and perceptions of comfort. To account for this, we obtained average outdoor temperature corresponding to the time period covered by the experiment for the 16th Arrondissement of Paris (in which the OECD is located) from the Weather Underground website (www.wunderground.com). Figure 3 plots these data, along with the average indoor temperature recorded by thermostats in our sample offices. The cold snap is obvious in this figure. Below we estimate regressions to estimate how this weather affected chosen temperature settings.

In addition to environmental factors, behavioral heterogeneity is important for understanding whether the ATEs we observe arise from an even, moderate response across the population or rather from a large response by a small portion of the population. To account for exogenous behavioural heterogeneity, we created an index which varied across offices (and not across time) as follows: For the baseline data—prior to treatment—we computed the average number of times that occupants of each office adjusted their thermostat by more than 0.25° up or down; this provides an exogenous indicator of how apt office occupants were to adjust their thermostats prior to the experiment. We call this new variable *Propensity*. Note that this variable only varies across offices, not time. Furthermore, it can be written as a linear combination of office fixed effects, and therefore in all regression models which include office fixed effects we can only include this *Propensity* variable by interacting with something that varies over time. Of course, we are interested in how the behavior of office occupants mediates the effects of defaults, and so we interact *Propensity* with the default setting: Our hypothesis is that office occupants who are more apt to adjust their thermostats are less susceptible to defaults, meaning that the regression coefficient on this interaction would be negative.

Indoor and outdoor temperature trends during the experiment.

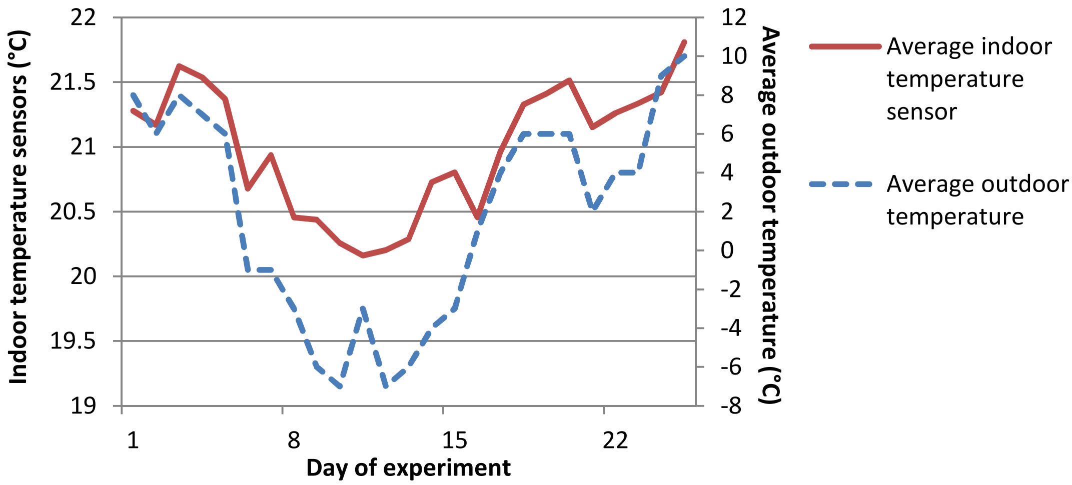

Finally, we hypothesize that social dynamics among office-mates may lead occupants of multi-occupant offices to "behave" differently than those who do not share an office. To analyse this hypothesis, we simply include in regressions a binary variable for whether the office has more than one occupant. The hypothesis to test is whether the regression coefficient on this variable is different from zero.

To analyse the impact of all of these factors, we focus on the marginal effects of the default setting on occupants' chosen temperature setting, as shown in the following regression:

$$s_{ih} = \beta_0 + \beta_D \text{default}_{ih} + \boldsymbol{\beta}_x \boldsymbol{x}_{ih} + \varepsilon_{ih} \tag{4}$$

where default_{ih} is the default setting for office i in hour h, and β_D is the effect of interest: the marginal effect of the default setting on chosen temperature setting s_{it}. The vector $\boldsymbol{x}_{ih}$ is a set of control variables, including weather effects and measures of behavioral heterogeneity, and $\boldsymbol{\beta}_x$ is the corresponding vector of regression coefficients to be estimated.

If we estimate equation (4) with standard linear regression techniques, our results would likely be biased by the fact that s_{ih} is constrained to be less than or equal to 3 degrees above the default setting, for reasons discussed in the previous section. To address this, we estimate Tobit regression models censored from above using the upper limits $L_{ih} \equiv \text{default}_{ih} + 3$ (Wooldridge, 2010).

Table 3 presents regression estimates for equation (4) using Tobit models with office fixed effects. Note that an ordinary least squares model along with uncensored random and fixed effects models are shown in this table for comparison. Using this table, we address, in turn, the effects of (a) weather, (b) heterogeneous preferences and perceptions, and (c) potential intra-office social dynamics on occupants' choices.

Table 3: Estimates from regressions of thermostat settings on default and other covariates.

Variables	*OLS*	*Office fixed effects*	*Office random effects*	*Tobit with office fixed effects*		
				(a)	*(b)*	*(c)*
Default	0.450***	0.385***	0.385***	0.251**	0.447**	0.453**
	(0.0872)	(0.0910)	(0.0906)	(0.125)	(0.188)	(0.192)
Avg. outdoor temp. (°C)	-0.0159	-0.0170	-0.0170	-0.0351**	-0.0344**	-0.0344**
	(0.0116)	(0.0108)	(0.0108)	(0.0143)	(0.0142)	(0.0141)
Default × Propensity					-2.420*	-2.419*
					(1.312)	(1.314)
Binary (Occupants >1)						0.184
						(4.729)
Default × (Occupants > 1)						-0.00881
						(0.236)
Observations	26,712	26,712	26,712	26,712	26,712	26,712
R^2 (generalized)[1]	0.070	0.060	0.070	0.080	0.084	0.084
Log-likelihood	-55,956	-45,831	-46,128	-38,882	-38,819	-38,819
Number of Offices	87	87	87	87	87	87

*Robust standard errors in parentheses, with *** p-value<0.01, ** p-value<0.05, * p<0.1.*
[1] The generalized $R^2 = 1 - (\bar{L}_0/\bar{L}_h)^2$, where $\bar{L}_h$ is the geometric mean likelihood in the sample for model h, and $\bar{L}_0$ is the corresponding quantity for an appropriately specified null model (e.g. including fixed effects). The generalized R^2 is equal the classical R^2 when both are computable.

Column (a) in Table 3 shows a Tobit model accounting for the effects of average outdoor temperature. This coefficient on weather is significantly negative, and indicates that a 1° decrease in outdoor temperature activates a 0.3° increase in occupants' thermostat setting. As with the DID estimates in Table 2, we see that the effect of defaults is still significant, even when accounting for data censoring: A decrease in the default setting resulted on average in a 0.25° decrease in the temperature setting that would have been selected had occupants not faced an upper limit on their setting. As expected, if we ignore censoring (the OLS, random effects, and fixed effects columns in Table 3), then the estimated effect of the default is much larger in magnitude while that of outdoor temperature is much smaller: The constraint forces those occupants with a high preferred temperature to lower their setting, and limits their ability to respond to colder weather with increases in their thermostats.

Column (b) in Table 3 shows the estimated impact of accounting for exogenous heterogeneous behaviour across offices. The regression coefficient on the Default×Propensity interaction variable is negative, as hypothesized. It is significant only at the 10% level (P-value=0.065). However, if we perform a test of whether this coefficient is strictly negative (as hypothesized above), then we reject with 3.25% significance the **null** hypothesis that occupants of offices who more frequently adjusted their thermostats *a priori* were **not** less influenced by the defaults. Figure 4 plots the estimated effects of defaults by how apt office occupants are to change their thermostats (using percentiles of the propensity index discussed above). We can see that the effect of defaults dissipates for offices above the 70th percentile of this propensity index.

Figure 3: Differential responses to defaults.

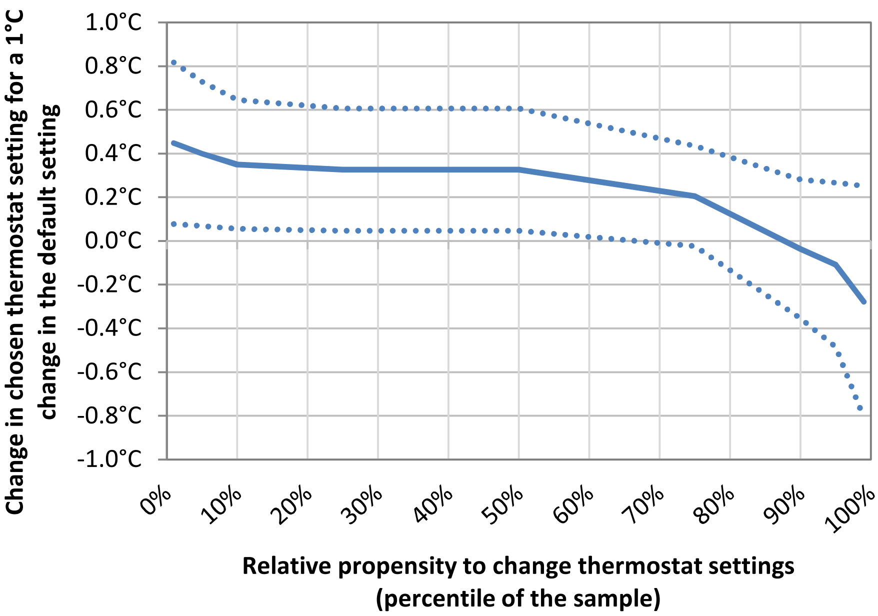

Note: Dotted lines represent a 95% confidence interval of the estimated effects.

Finally, column (c) in Table 3 shows that offices which have multiple occupants do not exhibit different patterns in their thermostat settings than offices with only one occupant. That is, we find no evidence that intra-office social dynamics play a role in choosing thermostat settings.

Discussion

The direct policy message from this small experiment is limited in scope, but clear: Small reductions in the defaults of office thermostats can lead to lower temperature settings by occupants in the winter heating season, which when scaled up to the whole building should translate into lower energy use. However, if the reduction in default temperature is too large, then occupants respond actively, increase their temperature settings, over-riding the effects of the change in the default setting. In quantitative terms, our results indicate that a reduction of the default temperature from 20°C to 19°C would decrease energy use, but a reduction to 17°C would have no effect.

There are a number of competing hypotheses which can explain our results here, but testing these hypotheses will require additional experimentation. The first hypothesis is motivated by signal detection theory (Green & Swets, 1966) and goes as follows: Occupants did not perceive small changes in their office temperature, but did perceive large changes, at which point they acted to improve their comfort by increasing their thermostat setting. Alternatively, small (and large) changes in office temperature were perceived, but acting to change the thermostat setting required cognitive and physical effort, the "cost" of which did not outweigh occupants' perceived gains from action. Lastly, changes in defaults may have reshaped occupants' temperature preferences and literally expanded their comfort zones, but when the default extended too far beyond what occupants were familiar with, then this comfort zone could no longer be expanded.

Identifying the correct explanation for our results is important for policies which would consider an intervention of the type described here, in which a default setting was manipulated so as to provoke more energy efficient choices among agents. If the signal detection theory explanation holds—and such an intervention capitalizes on agents' imperfect perceptions of temperature—then some might call into question the ethics of such interventions, which might be viewed as paternalistic (Mitchell, 2005; Sunstein & Thaler, 2003; Thaler & Sunstein, 2003). Even so, one could reasonably counter that many such default settings must be specified anyways (by building managers, policy makers, *etc.*), in which case why not select the one that is most energy efficient?

As behavioural interventions become more popular in public policy toolboxes (Abrahamse et al. 2005), the individual welfare implications of these interventions will become a more frequent topic of discussion. Consequently, it will be important to understand in more detail the mechanisms underlying individuals' responses to these interventions. This experiment has highlighted a simple, unambiguous way to lead agents to make more energy-efficient choices regarding their thermostat settings.

What are the policy implications of the results? It is important to emphasise that such measures cannot be considered as substitutes for more direct measures such as price-based instruments and direct forms of regulation. However, they can be a useful complement. In the case under analysis this is likely to be particularly true in the presence of split incentives or agency control problems which characterise a large number of relevant agents (tenants, office workers, etc). On the one hand, the incentives of a price-based measure will not be 'transmitted' to those undertaking decisions in terms of actual consumption. On the other hand, sole reliance on a direct form of regulation may constrain their capacity to make decisions which reflect their actual temperature preferences. The use of a behavioural 'prompt' alongside other measures - in this case a default setting – can result in reduced energy consumption in a manner which is cost-effective and consistent with utility maximisation.

REFERENCES

Abrahamse, W., Steg, L., Vlek, C., "Rothengatter, T. (2005). A review of intervention studies aimed at household energy conservation." *Journal of Environmental Psychology*, 25(3), 273-291.

Green, D. M., Swets, J. A. (1966). *Signal Detection Theory and Psychophysics*. Wiley, New York.

Johnson, E. J., Bellman, S., Lohse, G. L. (2002). *Defaults, Framing and Privacy: Why Opting In-Opting Out. Marketing Letters*, 13(1), 5-15.

Johnson, E. J., Goldstein, D. (2003). "Do Defaults Save Lives?" *Science*, 302(5649), 1338-1339.

Kahneman, D. (2003). "Maps of Bounded Rationality: Psychology for Behavioral Economics." *American Economic Review*, 93(5), 1449-1475.

Löfgren, Å., Martinsson, P., Hennlock, M., Sterner, T. (2012). "Are experienced people affected by a pre-set default option—Results from a field experiment." *Journal of Environmental Economics and Management*, 63(1), 66-72.

Madrian, B. C., Shea, D. F. (2001). "The Power of Suggestion: Inertia in 401(k) Participation and Savings Behavior." *Quarterly Journal of Economics*, 116(4), 1149-1187.

Marzilli Ericson, K. M., Fuster, A. (2011). "Expectations as Endowments: Evidence on Reference-Dependent Preferences from Exchange and Valuation Experiments." *Quarterly Journal of Economics*, 126(4), 1879-1907.

McKenzie, C. R. M., Liersch, M. J., Finkelstein, S. R. (2006). "Recommendations Implicit in Policy Defaults." *Psychological Science*, 17(5), 414-420.

Mitchell, G. (2005). "Libertarian paternalism is an oxymoron." *Northwestern University Law Review* 99(3), 1245-1277.

Pichert, D., Katsikopoulos, K. V. (2008). "Green defaults: Information presentation and pro-environmental behaviour." *Journal of Environmental Psychology*, 28(1), 63-73.

Sunstein, C. R., Thaler, R. (2003). "Libertarian paternalism is not an oxymoron." *University of Chicago Law Review*, 70(4).

Thaler, R. H., Sunstein, C. R. (2003). "Libertarian Paternalism." *American Economic Review*, 93(2), 175-179.

Wooldridge, J. M. (2010). *Econometric Analysis of Cross Section and Panel Data*, second ed. The MIT Press, Massachussetts.