THE
BEHAVIOURAL
INSIGHTS TEAM

# Evaluating the Nest Learning Thermostat

## Four field experiments evaluating the energy saving potential of Nest's Smart Heating Control

*A report by the Behavioural Insights Team*

*31 October 2017*


| | |
|---|---|
| Project lead & author: | Toby Park, Senior Advisor, Head of Energy and Sustainability |
| Project team: | Pantelis Solomon, Senior Research Advisor |
| | Michael Sanders, Chief Scientist and Head of Research |
| | Marius Dietsch. Associate Research Advisor |
| | Julie Faller, Research Advisor |
| | Daniel Gibbons, Research Advisor |
| | Elisabeth Costa, Director of Energy, Growth and Consumers |
| | Alex Tupper, Research Advisor |


**Structure of this report:**



**A long-form version of this report providing full details of all four studies has also been produced – please contact The Behavioural Insights Team for more detail.**

**Abstract**

Consumers often fail to set the most efficient schedules on their heating controls, leading to wasted energy and money. Smart heating controls employing sensors and machine learning algorithms may offer an attractive solution, and are therefore of key policy interest to the UK government. However, there is currently a lack of robust UK-specific evidence of their energy saving potential. To fill this evidence gap we present the results of four studies evaluating the Nest Learning Thermostat (NLT): 1. A propensity-score matched, quasi experimental study comparing 2248 self-selected NLT homes to a matched control group; 2. A small-scale 276-home randomised controlled trial (RCT) comparing the NLT to a 'modern suite' of controls (programmable timers, room thermostats and thermostatic radiator valves); 3. An exploratory, quasi-experimental study in which we run multiple analytical specifications on multiple samples (from a few hundred to a few thousand households) looking to test the robustness of savings against different analysis; and 4. An independent evaluation of data collected by Nest during a quasi-experimental, matched-samples study with approximately 20,000 homes, evaluating the impact of an additional opt-in algorithm called 'Seasonal Savings'.

Each of the studies presented challenges, most notably a self-selected treatment group and poor quality data in studies 1 and 3, and a relatively small sample (though greater methodological rigour) in study 2. Nonetheless, synthesizing the evidence across studies 1-3 we estimate that the 'basic functionality' of the NLT saves approximately 6-7% ($p<0.01$, 95% CI ±5.1%) of annual heating-system gas use (equating to 4.5-5% ($p<0.01$, 95% CI ±4.2%) of the total household gas consumption. These savings are relative to the 'full modern suite' of controls, and thus we conject that savings may be greater where replacing more rudimentary controls, though more research is needed to quantify this. We also find evidence that the NLT achieves these savings without undermining subjective thermal comfort. In study 4, running an intention-to-treat analysis, the opt-in Seasonal Savings algorithm is found to save 3.8% ($p<0.01$, 95% CI ±1.0%) of heating-system gas use across the whole targeted population (2.9% of total household gas use). 85% of eligible households opted in, meaning the average saving among those who adopted the algorithm, assuming only adopters achieved savings, is estimated to be 4.5% (95% CI ±1.2%) of heating system gas use (3.3% of total household gas use).

We further estimate that a NLT with Seasonal Savings is likely to have a return on investment of 4.5-6.5 years for medium-larger homes. This timeframe may significantly reduce if 1. More rudimentary controls than those used as our counterfactual are being placed, and/or 2. The NLT is self-installed or cheaper products with the same functionality become available.

Some important caveats remain in all cases and are discussed further in this report. Nonetheless these results represent the most robust studies of smart heating controls in the UK to date and provide a compelling case to support smart heating controls, where savings are empirically demonstrable, through UK government policy.

# Executive Summary

## Introduction

The heating controls in many UK households are outdated and difficult to use, with behavioural factors often leading to suboptimal heating schedules. For example, heating is often left on when homes are empty or kept unnecessarily high at night. Moreover, providing householders with advice is often not enough to reduce gas consumption.[1] Smart heating controls have the potential to reduce gas usage, lower carbon emissions, and save households money by automating and simplifying the user-experience, and using sensors and machine learning to encourage users toward more efficient heating schedules.

However, smart heating controls are a new category of technology and there remains a need for robust, independent and UK-relevant evidence on their energy-saving potential.  Though some estimates of savings do exist for certain' smart heating controls, these are mostly based on simulations and models. By contrast, this report aims to provide increased rigour and accuracy through field-trial evidence of real-world energy and cost savings. Such evidence can inform public understanding and is intended to be a predicate for evidence-based policy development.

The Behavioural Insights Team (BIT) has been commissioned by Nest, in collaboration with Npower, to fill this evidence gap. Acting as independent evaluators we undertook three field studies designed and evaluated by BIT, between 2014 and 2017, and one independent analysis of a study run by Nest. These studies were each designed to quantify the real-world energy savings of the Nest Learning Thermostat (NLT) to the highest standard of rigour possible given the available data and resource constraints. These studies are the first of their kind in the UK and we believe set the benchmark in terms of rigorous evaluation of smart heating controls.

[1] Department of Energy and Climate Change (2014).  Advice on how to use heating controls: Evaluation of a trial in Newcastle.

- **Study 1**: A quasi-experimental study comparing 2248 NLT owners against 2248 propensity-score matched non-NLT households, measuring dumb-meter gas billing data (2014-2015).
- **Study 2**: A randomised controlled trial (RCT) comparing 138 smart-metered homes with NLTs to 138 smart-metered homes with programmers, room thermostats and TRVs (the standard modern suite of heating controls), measuring smart-meter gas billing data and customer's subjective thermal comfort. (2016-2017).
- **Study 3**: A quasi-experimental study exploring multiple analyses and samples, each comparing several hundred NLT households to either several hundred (1-to-1 matching) or a few thousand (1-to-many matching) matched non-NLT households, measuring dumb-meter gas billing data. (2016-2017).
- **Study 4**: An independent evaluation of data collected by Nest during a quasi-experimental study comparing ~10,000 homes with NLTs, offered an opt-in Seasonal Savings feature, to a matched group of homes with NLTs, not offered Seasonal Savings, measuring heating schedule and boiler runtime data directly from the NLTs. (2016-2017).

### What is the Nest Learning Thermostat (NLT)?

The NLT is a smart heating control which allows intuitive manual control on the device and by mobile app, but which also learns your household occupancy, your heating preferences and the energy performance of your home and heating system. It aims to automatically set the most efficient schedule while maintaining user comfort.

In addition to sophisticated learning features, it also has weather compensation, and nudges users to choose more efficient heating settings by encouraging the collection of 'green leaves'. Users also receive Nest Home Reports via monthly email.

**Seasonal Savings** is an additional opt-in feature offered through Nest's Energy Partners which aims to find additional savings during the winter heating period, for example reducing temperatures by fractions of a degree overnight and other moments the user is deemed unlikely to notice. Any interaction on behalf of the user stops the incremental reductions.

## Main findings

**Note that throughout this report we express gas savings as:**

1. **Household gas savings – the reduction in total household gas consumption**
2. **Heating system gas savings – the reduction in gas used by the heating system (which typically accounts for 72-76% of household gas consumption)**[2]

- The NLT's basic functionality is estimated to save, on average, approximately **6-7% of the gas consumption of central heating systems**.
- This equates to a saving of approximately **4.5-5% of total annual household gas consumption.**
- This equates to savings in the region of **£25-27/year for medium homes**[3], and **£35-40/year for larger homes.**[4]
- Carbon savings are in the region of 110 to 160kg $CO_2$e per year for medium to larger homes.[5] If every UK home in receipt of an Energy Company Obligation (ECO) measure between 2013 and 2017 had received a Nest Learning Thermostat, this would equate to around 4.4-6.4 Mt$CO_2$e over the product lifespan[6].
- These savings figures are relative to a 'full suite' of modern controls - a programmable timer, room thermostat and thermostatic radiator valves. Though not quantified, we conject that savings may be greater when

---

[2] Using smart meter data from our Study 2 we estimate that approx. 76% of annual household gas consumption is used for space heating, and 83% over the winter period October-April. Ofgem's estimate for the UK average is 72%, annually.
[3] Assuming approx. 12,500kWh/year gas consumption (based on Ofgem data for median gas use), at 4.35p per kWh, Npower's current standard rate.
[4] Assuming approx. 18,000 kWh/year (based on Ofgem data for 75$^{th}$-percentile gas use), at 4.35p per kWh, Npower's current standard rate.
[5] Using 2017 gross calorific values for natural gas used in the UK – available from https://www.gov.uk/government/publications/greenhouse-gas-reporting-conversion-factors-2017
[6] 2 million homes have received an ECO measure since 2013. Assumed product lifespan 20 years.

compared to more rudimentary heating controls (e.g. 23% of UK homes currently have no thermostat).

- These savings are achieved with no loss of subjective thermal comfort, and we find weak evidence that thermal comfort may be being improved by the NLT[7].
- **NLT's opt-in Seasonal Savings algorithm achieves additional average savings of 4.5% of heating system gas use, among those who choose to opt-in. This equates to an additional £18-27 per year savings for medium-larger homes**. 85% of eligible homes do opt in, meaning the average savings across the whole targeted population (the intention-to-treat estimate) is 3.8% of heating system gas use. Expressed in terms of total household gas savings, these figures are 3.3% and 2.8% respectively.

**The above savings figures represent our best estimates based on the four studies undertaken. Note, however, that the NLT's 'basic functionality' estimates, derived from studies 1-3, have relatively wide confidence intervals reflecting some uncertainty. The magnitude of savings associated with Seasonal Savings has been estimated with greater certainty, reflected by narrower confidence intervals. See pages 11-13 for more detail.**

## Policy implications

This research has been policy-focused from the outset and has aimed to address the major questions raised by UK policymakers over the course of its design and implementation. Given the rigour of these studies and the demonstrated energy savings, we believe that this evidence provides a strong case to incorporate smart learning heating controls into existing and future energy efficiency policy. At the time of publication, this is particularly pertinent given the UK Government's recently published Clean Growth Strategy, alongside a consultation on the UK's Green Deal,[8]

---

[7] The results of our subjective thermal comfort analysis based on occupant surveys are limited by a small sample and some risk of bias. Nonetheless, the evidence is strong enough to rule out, with a reasonable degree of certainty, the possibility that the NLT is *undermining* comfort.

[8] The 'Green Deal' is the UK government's flagship policy aiming to encourage and support able-to-pay households to install energy efficiency retrofits by providing financial support in a pay-as-you-save scheme.

as well as recent announcements of more stringent boiler regulations (including the mandating of additional efficiency measures to be installed with new boilers, including heating controls).

However, incorporating this evidence in UK policy may not be straightforward: most existing energy policies (including the Energy Company Obligation (ECO), Energy Performance Certificates (EPCs), building regulations and The Green Deal) currently base their energy-savings estimates on an engineering model of building efficiency – the Standard Assessment Procedure (SAP). This approach has two major implications:

- SAP is "*based on standardised assumptions for occupancy and behaviour*", assessing "*how much energy a dwelling will consume, when delivering a defined level of service provision*".[9] In other words, it assumes the building requires a certain amount of heat for a certain number of hours per day, and calculates the energy needed to achieve that - a well-insulated home will use less, for example. However the NLT aims to save energy by *reducing* the heating hours and the temperature at moments that are acceptable to the occupant, such as when the home is empty. This mechanism of savings undermines SAP's core assumptions, and as such SAP (and the energy policies built upon it) cannot readily recognise technologies such as the NLT. Therefore, if new policy measures are to support smart heating controls, they will need to be designed in a way which circumvents SAP, or SAP will need to be fundamentally altered to accommodate a more sophisticated model of energy use incorporating dynamic assumptions about the level of service an occupant needs.
- The SAP-derived energy savings estimates associated with traditional products such as insulation are not comparable to the savings estimates we present here for the NLT. SAP is known to produce unreliable and often very optimistic savings estimates[10], giving little perspective on real-world performance where

[9] Quotations from https://www.gov.uk/guidance/standard-assessment-procedure
[10] Kelly, S, Crawford-Brown, D & Pollit, M (2012), Tyndall Centre for Climate Change: Building performance evaluation and certification in the UK: is SAP fit for purpose?

human behaviour and imperfect installation can significantly diminish performance. In contrast, we present real-world achieved energy savings for the NLT.

We acknowledge that energy efficiency models such as SAP can play an important role in estimating the performance of the building fabric, separate to its use and occupancy. However this approach is becoming increasingly unfit for purpose where disruptive technologies and demand-side interventions put more onus on influencing users and their actual energy consumption rather than solely on physical properties of the building fabric and services. Placing greater weight on real-world field trial evidence such as that presented in this report is the obvious way to overcome these shortcomings, but is also fundamentally more accurate, where research is good quality (by prioritising real-world savings over modelled estimates), is likely to bolster innovation (by implicitly supporting a wider variety of energy efficiency measures), and can lead to UK consumers having more choice, and more ways to save energy and money.

## Cost Savings

Table 1 below summarises the savings and payback periods of the most common efficiency retrofits in the UK. Two key points can be taken from this table: 1. The SAP estimates in the left-hand portion of the table clearly overestimate the savings relative to their real-world counterparts in the right-hand portion, and 2. The NLT appears to be very competitive in its payback period compared to the real-world performance of other products, though it risks being overlooked if compared directly to the SAP-based estimates.

It is also worth highlighting that though not included in this table, the NLT is likely to perform better than indicated on two counts: first, as discussed below, these NLT savings are relative to the full 'modern suite' of controls, and we anticipate that savings are likely to be greater when replacing more rudimentary heating controls, potentially reducing the payback period to less than 3 years in larger homes (though further research is required to quantify this). Second, smart heating controls may be available for significantly less than the £279 installed cost assumed here – for example many users install their own NLTs (saving £60-80), and Nest have recently released their 'Nest Thermostat E', currently available only the in US, but offering the same functionality at a 33% price reduction. This would also significantly reduce the pay-back period.

| | SAP-based modelled savings estimates | | | Real-world field trial / quasi-experimental evidence | | | | | |
|---|---|---|---|---|---|---|---|---|---|
| | Loft insulation | Cavity wall insulation | External solid wall insulation | Condensing boiler | Loft insulation | Cavity wall insulation | External solid wall insulation | Basic NLT replacing thermostat, programmer & TRVs | NLT + Seasonal Savings replacing thermostat, programmer & TRVs |
| Medium home % gas saving | - | - | - | 8% | 3.8% | 7.15% | 13.2% | 4.7% | 8%[11] |
| Medium home £ savings | £120 | £95 | £160 | £44 | £21 | £39 | £72 | £26 | £46 |
| Medium home product cost | £285 | £423 | £8,000 | £2,000 | £225 | £423 | £8,000 | £279 | £279 |
| Medium home payback time | **2.5 yrs** | **4.5 yrs** | **50 yrs** | **46 yrs** | **11 yrs** | **11 yrs** | **112 yrs** | **11 yrs** | **6.5 yrs** |
| Larger home % gas saving | - | - | - | 8% | 3.8% | 10.5% | 13.2% | 4.7% | 8% |
| Larger home £ savings | £225 | £225 | £425 | £63 | £30 | £82 | £103 | £37 | £63 |
| Larger home product cost | £395 | £720 | £20,000 | £2,000 | £395 | £720 | £20,000 | £279 | £279 |
| Larger home payback time | **2 yrs** | **3 yrs** | **47 yrs** | **32 yrs** | **13 yrs** | **9 yrs** | **193.5 yrs** | **7.5 yrs** | **4.5 yrs** |

*Table 1: SAP-modelled and real-world energy savings estimates for energy efficiency products. All assumptions and sources in footnote[12]. All £ figures rounded to the nearest £. All payback periods rounded to the nearest 0.5 years.*

[11] To generate an estimate for the NLT with Seasonal Savings, we have summed the savings estimate of the basic NLT, and that of the Seasonal Savings Algorithm beyond the basic NLT. Though this is intuitively correct, and likely to be accurate within the context of the relatively wide confidence intervals of all findings in this report, we caveat this with the warning that the studies from which the two figures are based differed significantly in their methodology and their sample – see the technical summaries for more detail.

[12] Real-world boiler and insulation savings are from: Department of Energy and Climate Change (2013) National Energy Efficiency Data-Framework. Part II Impact of Energy Efficiency Measures in Homes. For cavity

## Savings compared to other heating controls

The savings figures for the NLT summarised above are compared to a full modern suite of controls: a programmable timer, a room thermostat and thermostatic radiator valves (TRVs). These are the most common heating controls in the UK, covering 49% of homes. However, many households have more rudimentary controls: 23% of UK homes do not have thermostats, 2% have no heating controls at all, whilst many have no TRVs, and some have only a thermostat, or radiator valves with a boiler on/off switch but no programmable timer[13].

We might expect the savings associated with installing a NLT to be greater when replacing more rudimentary controls, however we cannot accurately quantify this additional benefit due to a lack of real-world evidence of the relative performance of other heating controls.[14] For example a recent 2017 systematic review of UK and international evidence commissioned by the Department for Business, Energy and Industrial Strategy concluded that there are 2%-10% savings associated with TRVs

---

wall insulation (n=43,000 homes), data is available segmented by house size, and so we have used 'detached' for 'larger home', and the average of 'end terraced' and 'mid terrace' for 'medium home'. For other products, samples are smaller and so no segmentation by house size is available. SAP estimates are those provided by the Energy Savings Trust, and we use 'detached' for 'larger home' and the average of 'terraced' and 'semi-detached' for medium homes (deemed most similar to DECC's categories). For the SAP estimates, savings are only provided in £, not in % gas. To convert the other % savings to £ savings, we assume a unit rate of 4.35p/kWh (typical standard variable tariff, based on Npower's rates at the time of writing), and assume 'larger homes' to use 18,000kWh/year and 'medium homes' to use 12,500kWh/year, in-line with Ofgem's data 75th percentile and median consumption respectively. All insulation product costs are those provided by the Energy Savings Trust, again based on 'detached' (larger homes) and averaging 'terraced' and 'semi-detached' (medium homes). The NLT cost is the full installed RRP. The boiler cost is an average obtained from searching various providers, including installation, and is relatively conservative (low). Further information on typical boiler + installation costs can be found at https://www.theheatinghub.co.uk/guide-to-boiler-installation-costs.

[13] Department of Energy and Climate Change (2014) How heating controls affect domestic energy demand: A Rapid Evidence Assessment. (see page 25 for an overview of the penetration of different heating controls)

[14] DECC reference in their literature a large US study (RLW Analytics, 2007) finding 6% savings associated with 'modern programmable thermostats' (often called setback thermostats in the UK) compared to matched homes without such controls (the control group in this study would most likely have had manual thermostats without timers or TRVs). However setback thermostats are not the same as separate programmers, room thermostats and TRVs – in some ways more sophisticated but also often criticised for delivering poor savings due to being difficult to use and often left un-programmed. This evidence is also not UK-specific. In the UK, the Energy Saving Trust (EST) claim that installing and correctly using a programmer, room thermostat and TRVs would save a typical 3-bedroom semi-detached house £75 per year (equating to around 21% of heating system gas use, or around 15% of household gas use). However, this is likely to be an overestimate, being based on SAP and thus not accounting for behavioural and other in-use factors which might undermine savings.

(but with very low confidence), but concluded there to be a 'lack of robust evidence' for programmable timers and room thermostats, as well as for most other heating control types[15].

It is important to address this evidence gap, as the optimal policy and marketing response would be to target NLTs at homes with the least sophisticated controls where savings may be greatest. We therefore suggest further work is necessary to measure the savings of (and infer the NLT's savings relative to) other types of heating controls. If reliable evidence does not exist, we suggest seeking guidance from BEIS or the Building Research Establishment (BRE) on a safe working estimate to fill this evidence gap and help to inform policy development. Figure 1 below summarises the 'stacked' savings and the location of this uncertainty.

[15] Oliveira, S et al (2017) Heating Controls: International Evidence Base and Policy Experiences. University of the West of England and BSRIA, prepared for the UK Department for Business, Energy and Industrial Strategy.

*Figure 1: Incremental 'stacked' savings estimates of the NLT and Seasonal Savings compared to other heating control types*

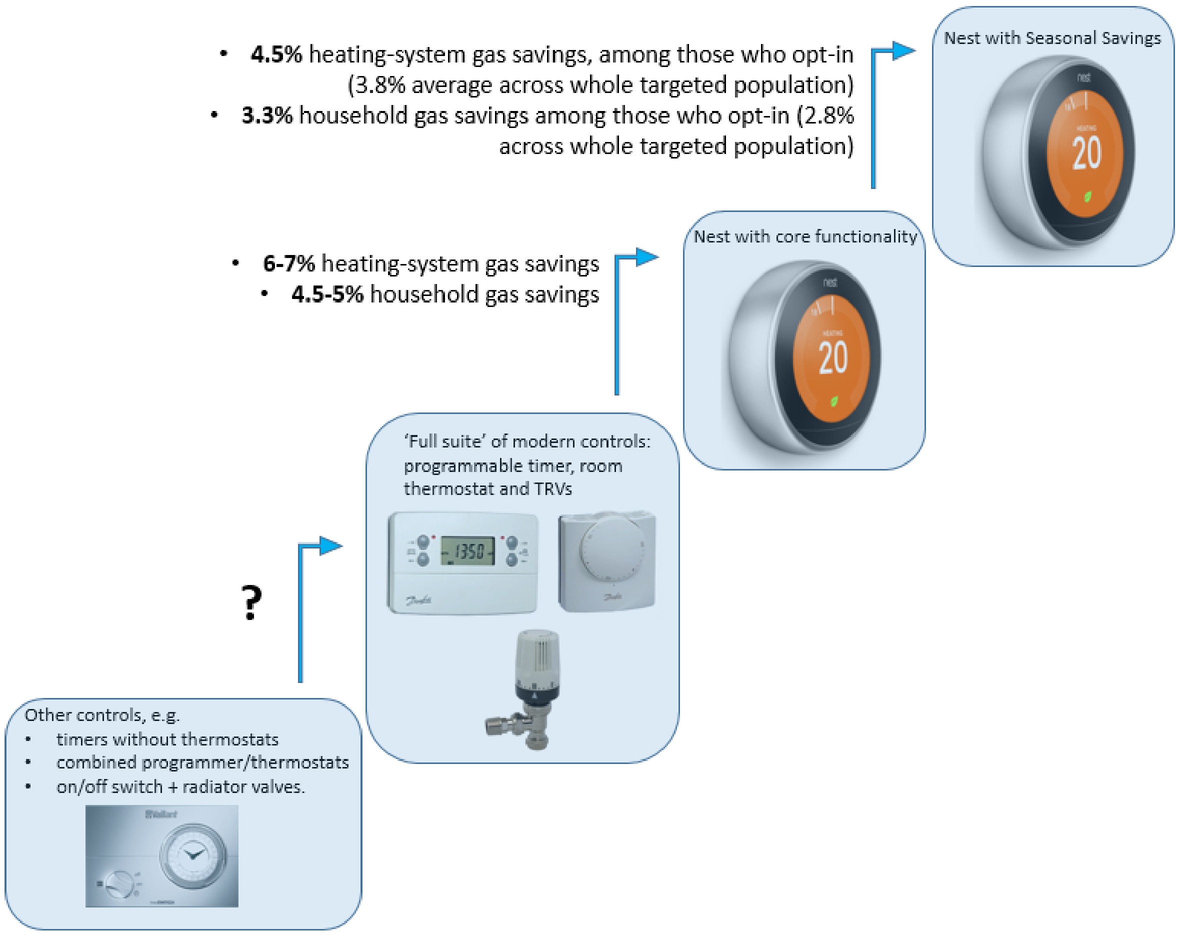


## Conclusions

These four studies conducted by BIT provide evidence that the NLT saves energy in the typical UK context of gas central heating. Moreover, these savings are not trivial, likely to be in the region of 6-7% of heating-system gas use (4.5-5% of annual household gas use), for the basic NLT functionality compared to households with programmable timers, room thermostats and TRVs. We are not able to estimate the savings compared to more rudimentary controls, though we conject that they are likely to be greater, and maybe significantly so. Moreover, we estimate that the Seasonal Savings algorithm saves an additional 4.5% of heating system gas use (3.3% of total household gas use) among those who adopt it. With 85% of eligible customers choosing to adopt it, averaged across a targeted population this equates to an additional 3.8% of heating system gas use (2.8% of household gas use).

Each individual study displays a degree of uncertainty around the precise magnitude of savings, reflected by the confidence intervals which are wide for studies 1-3 and relatively narrow for study 4. Each study also had further important caveats, outlined in the technical summaries below and detailed in the full report.

These caveats accepted, the relative consistency of savings evidenced across the studies adds confidence to our conclusion that the NLT provides substantial energy savings, and we believe this evidence builds a strong case for either developing new policies which support smart heating technologies, or taking a renewed look at the role of SAP to identify ways in which existing policies can be modified to accommodate such technologies.

Moreover, the methodologies detailed in this report set the benchmark in robust evaluation of real-world energy-savings in the UK. With behaviour-change, demand-side, and smart technology interventions becoming increasingly commonplace, laboratory tests and modelled energy savings encompassing static assumptions about occupants' energy needs will become increasingly inadequate. Their shortcomings can be addressed by supplementing these conventional techniques with field trial methodologies similar to those used in this report. Aligned with this agenda, the BIT continues to support the Department of Business, Energy & Industrial Strategy, Ofgem and others in developing field trial evidence across a range of energy, sustainability and consumer issues.

# Technical Summary

This section provides a short summary of the four studies' methodologies, results and caveats. Each study has been designed to be as rigorous as possible within resource and practical constraints: each uses a slightly different approach, aiming to make best use of the data sources available in each case, while simultaneously maximizing external and internal validity and addressing the most important research questions.

## Study 1 – Technical summary

### Study 1 methodology

Study 1 was a quasi-experimental matched design making use of 2248 existing NLT customers and 2248 non-NLT customers, propensity-score matched on historical consumption (normalised for weather variation using Heating Degree Days (HDDs)) and a range of household and postcode level characteristics. We drew upon dumb-meter gas consumption data (billing data) from 2009 to 2015, excluding estimated readings, and ran a regression on this matched sample using a fixed effects model (an extension of a standard difference in differences model for panel data with more than two time periods). This allows us to observe the trends in energy consumption of households who have opted into the NLT, and compare them to a matched sample of non-NLT households, whilst controlling for observable characteristics. We also run a number of balance checks, representativeness checks and robustness checks.

**Study 1 findings**

- **Annual household gas saving – 5.8%** ($p<0.01$, 95% CI ±3.2%), compared to aggregate of other heating controls.
- **Heating-system gas saving – 7.8%** ($p<0.01$, 95% CI ±4.3%), compared to aggregate of other heating controls.[16]
- **Estimated cost savings of £30-35/year (medium homes) - £40-50/year (larger homes)**[17]

In Study 1 we find an estimated 5.8% ($p<0.01$, 95% CI ±3.2%) reduction in household gas consumption due to the NLT. This result is highly statistically significant ($p<0.001$). Averaged across the year, this equates to around 2kWh per day for an average 'medium' household (typically using 12,500kWh/year). This equates to a roughly £30-45 annual saving in gas bills.[18]

We also estimate the savings on the heating system. Drawing upon Study 2's smart meter data, we estimate that heating is typically responsible for 76% of annual household gas consumption, while the UK government estimates it to be 72%[19]. We use the mid-point (74%) here and in all other similar conversions in this report, and thus infer that the NLT saves approximately 7.8% (±4.3%) of heating system gas use.[20]

Note that these results are annual savings figures (some results from other studies are winter-only) and that some users in Study 1 had the optional Seasonal Savings feature enabled (we isolate the savings of this feature in Study 4, and of the NLT without this feature in Study 2).

[16] Assuming 74% of household gas consumption is for heating, as per Ofgem's data (72%) and our data from Study 2 (76%). Thus, 5.8%/0.74 = 7.8%
[17] Medium homes assumed to use 12,500kWh/year (Ofgem's 'median' household), and larger homes 18,000kWh.year (Ofgem's 75th percentile). Assumed cost 4.35p/kWh based on Npower's standard variable tariff.
[18] Based on a unit cost of 4.35p, for a household using 12,500-18,000 kWh/year, which are the 'medium' and 'high' consumption thresholds stated by Ofgem.
[19] Data available from http://webarchive.nationalarchives.gov.uk/+/http:/www.berr.gov.uk/files/file43843.pdf
[20] Simply, 5.8%/0.74 = 7.8%

**Study 1 conclusions and caveats**

We conclude from Study 1 with a reasonable degree of confidence that the NLT saves energy, and that savings are unlikely to be trivial, though urge some caution in interpreting the precise magnitude of savings since the confidence intervals around the point estimate are relatively wide (±3.2%), and the counterfactual is not well identified, since some NLT households will have had the Seasonal Savings feature enabled (we isolate this feature in Study 4), and the control households has an unknown assortment of heating controls.

We also caveat this study due to its reliance upon a self-selected sample. Though we produce a good match between the NLT group and the control group, we can only match on characteristics we have data on, and the low predictive power of the propensity score suggests that the groups may not be well matched on other, unobserved, characteristics which predict the purchasing of a NLT (attitudes towards new technology, for example). This is not necessarily a problem, since any imbalance between groups which impacts gas consumption would, if consistent over time, be controlled for by proxy since we control for historic gas consumption. However we cannot rule out the possibility of imbalance on time-variant characteristics, for example, it is possible that households which have purchased a NLT would also be motivated to make other contemporaneous energy-saving changes in their behaviour or their homes, confounding the savings estimate. Nonetheless, despite these caveats, at the time of completion in 2015 Study 1 represented the most rigorous evaluation of its kind in the UK, and presents evidence of the NLT's energy saving potential.

## Study 2 – Technical summary

**Study 2 methodology**

Though providing promising evidence, Study 1 was exploratory and was always intended to be followed by a more rigorous randomised controlled trial (RCT). Study 2, run over the winter of 2016/17, therefore aimed to achieve three things:

- Improve on the methodological limitations of Study 1 by running a RCT. This aims to remove self-selection bias and provide improved causal attribution by randomly selecting half of the households to receive a NLT.
- Harness newly available smart meter data, which, compared to dumb-meter data used in Study 1, provides more accurate, less noisy, more reliable, and more frequent gas consumption data. This significantly increases our precision and statistical power, allowing us to run a relatively smaller sample size. This smaller sample is in any case necessary for a RCT, given the expense of providing free NLTs to a random population.
- Address two key unanswered question identified as important to policy colleagues in the Department of Business, Energy and Industrial Strategy (BEIS), formally the Department of Energy and Climate Change (DECC). Namely:
    - *What is the counterfactual against which the NLT can claim to save energy?* In study 1 we demonstrate an estimated 5.8% gas saving, compared to an unknown aggregate of heating technologies in the control group. In Study 2 we therefore compare the NLT specifically to a 'full modern suite' of technology, defined as a programmable timer, a wall thermostat, and thermostatic radiator valves[21]. All control households, and all NLT households before given a NLT, have these standard heating controls, as recommended by BEIS as a well understood, and very common benchmark.
    - *Can the NLT save energy without compromising occupant comfort?* Unlike most traditional energy efficiency products, such as insulation, the NLT saves energy by reducing the amount of heat delivered to the

[21] According to the Department of Business, Energy and Industrial Strategy (BEIS), timers have 97% penetration, wall thermostats have 77%, and TRVs 66%. 49% of UK households have all three. This counterfactual was chosen because 1. It is the most common combination in the UK, 2. BEIS colleagues suggested to us it was the 'most well-understood' heating control technology, and 3. This is logically where the burden of proof lay, as the NLT has the functionality of a programmable timer and thermostat (and more), and was thus already presumed by BEIS colleagues to be as good as these conventional controls – the requirement on our research was therefore to demonstrate whether or not the NLT saves *more* than these conventional controls.

building. If such technology is ever to receive policy support, there is a clear need to demonstrate empirically that the NLT can achieve these savings without diminishing subjective comfort. We therefore measure subjective comfort directly through surveys in Study 2.

Our approach is therefore to run a randomised controlled trial comparing 139 randomly allocated NLT households to 139 control households. Since the sample is relatively small, we minimise the risk of imbalance between groups by running a matched-pair randomised design in which participants are paired to their closest match (on gas consumption and other household characteristics), with each pair randomly split into a treatment and control household. We also run balance checks on all observable characteristics, finding good balance at the outset, but some imbalance emerging after installations are complete due to significant attrition at this stage (some homes in the NLT group refused the NLT, while others were not contactable). We attempted to control for this imbalance by 1. Only including households, in both the NLT and control groups, from a pool of household all of whom expressed consent to be in the trial and receive a NLT, 2. Removing matched control homes when a NLT home drops out, and 3. Controlling for imbalanced covariates in the regression analysis. Despite these efforts, a degree of self-selection bias is unavoidably introduced since NLT households are free to refuse the NLT when installation is attempted, and it is impossible to perfectly mimic this attrition in the control group.

In Study 2 we exclusively draw on homes in Yorkshire and the Midlands to reduce weather variance and increase the practicality of installations within a narrow timeframe. This is also a region with high concentration of smart meters among Npower customers, allowing us to exclusively use smart meter data. Analysis utilises a panel dataset with frequent smart meter reads before and after NLT installation for both groups. This spans the winter heating period 1st October 2016 – 30th April 2017, with the NLT installations taking place in January and early February 2017. Though we only expect savings to occur in the winter heating period, as secondary analysis we also look at data from a wider period of observation, from June 2016 to June 2017.

We run regression analysis on the panel dataset, and apart from the treatment variables, our explanatory variables included the average number of heating degree days (HDDs) between two readings. Additionally, we used dummy variables for each month, which allowed our model to capture seasonal effects that are not captured by HDDs (seasonal effects exist beyond those driven by temperature, for example due to sunshine in the summer providing solar gain not captured by HDDs from weather stations, and other behavioural factors such as preferences for heating in the winter at temperatures which would not require heating in the summer). We also include demographic variables as further covariates to control for those characteristics known to be imbalanced. In all regression models we clustered our standard errors on the household level in order to take into account the fact that household's readings are correlated over time.

We also isolate the first three weeks of data after installation, using a second treatment variable, as the NLT takes time to learn before it can save energy. This three week period was identified as appropriate by analysing set-point temperature data from the NLT's themselves, comparing more recently installed to less recently installed NLTs to determine when they had finished learning, though for robustness checks we also test shorter and longer presumed learning periods. Note that due to the learning period, though most households in the NLT group had their installation in January and early February, the group is not truly 'treated' until mid-late February, giving us relatively limited post-installation data.

**Study 2 findings**

- **Winter household gas saving**, compared to a modern suite of heating controls: **5.6%** ($p<0.01$, 95% CI ±4.2%).
- **Annual household gas saving**, compared to a modern suite of heating controls:
  - Estimate 1: **4.5%** ($p>0.05$, 95% CI ±5.6%)[22]
  - Estimate 2: **5.0%** ($p<0.01$, 95% CI ±4.2%)[23]
- **Heating-system gas saving**, compared to a modern suite of heating controls:
  - Estimate 1: **6.4%** ($p>0.05$, 95% CI ±7.6%)[24]
  - Estimate 2: **6.7%** ($p<0.01$, 95% CI ±5.7%)[25]
- **No loss of subjective thermal comfort**
- **Estimates cost savings of £25-27/year (medium homes) - £35-40/year (larger homes)**[26]

Our primary analysis shows an average 2.88 kWh/day saved from a baseline of 51.46 kWh/day across the winter observation period February to April. This equates to a 5.6% ($p<0.01$, 95% CI ±4.2%) saving in household gas consumption. We also demonstrate that the average daily consumption over this period (51.46kWh/day) is almost identical to the average consumption over the full winter heating period October to mid-May (52.0kWh/day), and as such, we conclude that this savings estimate is likely to be a good estimate for the full winter heating period.

Our secondary analysis expands the observation period to include some warmer months, June to June, finding an average saving of 1.69kWh/day from the baseline of 37.56kWh/day. This equates to 4.5% (±5.6%) saving in household gas consumption. This finding is not statistically significant, so should be treated with caution, though we note that the result is in-line with expectations, since savings are only expected to occur in the winter heating period, and so this estimate is essentially a diluted version

[22] Inferred from our 'extended' observation period Feb-June, over which the average daily consumption was nearly identical to the average daily consumption of the whole year.
[23] Inferred from our winter observation period (primary result) and conservatively assuming that savings only occur during the winter period October – mid-May.
[24] Inferred from the annual savings estimate, on the basis that 74% of annual consumption is for heating.
[25] Inferred from the winter savings estimate, on the basis that 83% of winter consumption is for heating.
[26] Based on a 'middling' estimate of 4.75% annual savings, on 12,500kWh/year consumption (medium homes, as per Ofgem's median household data) – 18,000 kWh/year consumption (larger homes, as per Ofgem's 75th percentile household data). Assumed cost 4.35p/kWh as per Npower's standard variable tariff.

of the primary winter savings figure above (indeed this can be confirmed quantitatively, by dividing the 5.6% savings accrued in the shorter winter observation period by the total consumption across this longer observation period).

We also show that the average daily consumption over this Feb-June period (37.56 kWh) is very close to that for the full year (36.0kWh/day), and so conclude that this savings estimate is likely to be a good estimate of the full annual savings in household gas. As an alternative approach to estimating the annual savings, we conservatively assume savings only occur over winter, and so starting with our primary winter result of 5.6% or 2.88kWh saved per day, across the winter heating period October to mid-May, this equates to an annual saving of 5.0%. This slight discrepancy between our two annual savings estimates (4.5% and 5.0%) impossibly suggests that aggregate savings in the winter are slightly higher than aggregate savings across the whole year, highlighting the margins of error in all of these results. Nonetheless, these estimates are remarkably close, adding confidence to our conclusion that the annual savings are likely to be approximately 4.5-5% of household gas consumption.

This equates to bill savings of £25-27 per year for medium homes (defined at the median gas user according to Ofgem, who use 12,500 kWh/year) and £35-40 per year for larger homes (set at the 75th percentile gas user, who use 18,000kWh/year).[27]

We also convert these figures to an estimate of savings on the heating system by using the smart meter data to isolate the baseline (non-heating) gas consumption. This approach suggests that over the winter heating period, 83% of gas consumption is driven by heating. We therefore infer that the saving on heating system is 6.7%.[28] Across the full year, our data suggest than 76% of annual consumption is used for space heating, close to Ofgem estimates of 72%. Our estimate may be slightly high as we cannot account for the fact that winter bathing and cooking is likely to use slightly more gas than summer bathing and cooking. We therefore use a middling estimate of

[27] In all our cost saving estimates we assume a unit price of 4.35p/kWh, typical for a standard variable tariff, and in-line with Npower's standard tariff at the time of writing.
[28] Simply, 5.6%/0.83

74% throughout this report. This implies heating-system savings of 6.4%[29]. We therefore conclude that the NLT is likely to save in the region of 6-7% of heating system gas use.

Note that unlike Study 1, these savings are relative to the 'full modern suite' of heating technology of a programmable timer, wall thermostat and thermostatic radiator valves. We conjecture that the savings may be significantly higher when compared to more rudimentary heating controls, though more research is needed to corroborate this, as the relative performance of existing heating controls is poorly evidenced in the literature. Study 2 results also represent the NLT with core functionality only, with no homes having the Seasonal Savings feature enabled. As such, the savings identified within Study 4 are in addition to these estimates, though being a separate study with a different sample and different methodology, a simple addition of savings figures from Studies 2 and 4, though the best we can do with available evidence, may be prone to some error.

Figure 2 shows the gas consumption per heating degree day, highlighting the energy savings which largely emerge because the NLT group is exposed to fractionally colder weather after installation. Note, however, that Figure 2 does not reflect the final specification for our analysis, in which we control for weather by including heating degree days as a covariate, rather than as a denominator to gas consumption - we include this figure as it is the most intuitively interpretable graphical representation of weather-controlled gas consumption.

[29] Simply, 4.75%/0.74, using 4.75% as a middling estimate of our various estimates for the annual savings (4.5% and 5.0%)

*Figure 2: Gas consumption per heating degree day across winter heating period*

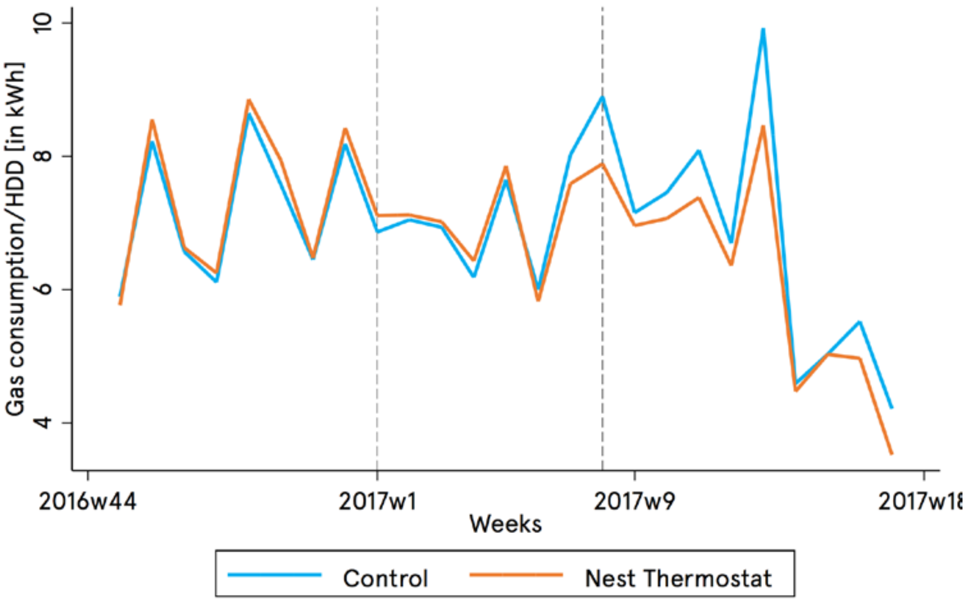


*Note:* The first vertical line denotes the week of the first and the second vertical line denotes the day of the last Nest installation.

We also run analysis on subjective thermal comfort, with data collected through two surveys: one before NLT installations in the mid-winter and one at the end of winter, each for both groups. This allows us to see how comfort has changed over the course of the winter, and the extent to which the NLT groups' comfort has changed relative to that of the control group after receiving a NLT. As might be expected, we see both group's comfort increase between December and April (since the weather is gradually warming), but the NLT group's comfort increases more. This difference is weakly significant (p=0.065). There are some problems with this analysis, however: not all participants answered both surveys and so the sample is relatively small (n=196, of a possible 278) and we see differential engagement with more of the NLT group (118) answering the second survey than the control group (78). This may introduce some bias, though we do not know how severe or in which direction. We therefore interpret this finding conservatively, and do not infer that the NLT improves comfort. However, we suggest the results are strong enough that we can conclude with some confidence that the NLT does not diminish comfort as a consequence of its energy saving.

**Study 2 conclusions and caveats**

Study 2 provides further strong evidence that the NLT saves energy, though again the point estimates are surrounded by relatively wide confidence intervals reflecting a fairly high degree of uncertainty about the precise magnitude of savings. However, the internal consistency of the results within Study 2, and with those from Study 1 (which used a very different sample, data source and methodology), and with Study 3 (which uses a different methodology again) is encouraging in this regard and adds a lot of confidence to our overall conclusion that the NLT saves energy of approximate magnitudes outlined above.

However, some caveats do remain. Principally, there was a high degree of attrition during installations, with many households in the NLT group either refusing installation, or being uncontactable (despite all participants in both NLT and control groups being selected from a pool of people who had consented). This not only reduces our intended sample size (from 400 to 278), but also introduces a risk of self-selection bias, since the remaining NLT households are now filtered toward certain pro-NLT attitudes. To minimise this risk, we also removed the matched-pair control household whenever a NLT household attrited. Nonetheless, we would expect some bias to emerge. This is indeed seen with the final sample showing some imbalance on attitudes to new technology (NLT group viewed technology slightly more favourably), income (NLT group slightly wealthier) and number of children (NLT group having slightly more children on average). We control for these variables in the analysis. Moreover, as per Study 1, we stress that imbalance between groups on variables such as these, which are fairly consistent over time, is not necessarily problematic, because if they impact gas consumption, then we control for them indirectly by controlling for historic gas consumption – however any imbalance somewhat undermines ability to claim causality by undermining the purity of the counterfactual.

Perhaps more importantly, some imbalance in geographic location also emerges as a result of the attrition, which leads to imbalance in weather which is *not* consistent over time - the NLT group are exposed to slightly colder weather, but predominantly after installation. The NLT group 'overcoming' this colder weather is a contributing driver of the savings estimate. Our analysis specification is well suited to account for this imbalance, provided the imbalance is 'real' – the risk of confound is from unreliable data imbalanced between the groups. Specifically, we cannot fully rule out the possibility that if one of our 12 weather stations (near which more NLT households are located that control households) is giving spuriously cold data over the second half of the winter, this could lead to the erroneous conclusion of gas savings (i.e. a Type 1 error). This seems reasonably unlikely, but is exactly the kind of confound which a perfect RCT without any self-selection bias or imbalance between groups aims to avoid. This imbalance is therefore an unfortunate and unavoidable consequence of installation refusals among the NLT group, though on balance and given the consistency of findings across studies and robustness checks, we nonetheless conclude this risk of confound is relatively minor.

## Study 3 – technical summary

### Study 3 methodology

Study 3 was a quasi-experimental study in which we took administrative gas readings, from dumb-meters, from Npower for approximately 4,200 NLT customers and 27,400 non-NLT customers, with data extending from 2009 to 2017.

In many ways Study 3 is similar to Study 1, being quasi-experimental, relying on dumb-meter data, and comparing self-selected NLT households to a matched cohort without NLTs, who have an unknown mixture of other heating controls. However, with more post-installation data available, Study 3 was an opportunity to replicate the findings and to attempt to improve on the methodology of Study 1. The main differences to Study 1 are to do with the way we created a 'pre-intervention' and 'post-intervention' timeframe. The NLT installation date is different for each household and spans multiple months. In Study 1, we dealt with this crudely by arbitrarily defining the first NLT installation date in our sample as the intervention date, though gas readings were very sparse and in any case did not align with the intervention date for many households, meaning the distinction between pre- and post-intervention data was fuzzy. This was unavoidable as NLTs were relatively new and thus relatively little post-installation data was available. With a longer observation period at our disposal in Study 3, and with NLTs having been installed in the UK for longer, we are able to define a precise pre-intervention period (January 2014 – February 2015), and a precise post-intervention period (December 2015 – March 2017). All NLT households included in the study had their NLT installed between these two periods, and we required at least two reliable (not estimated) meter reads in both the pre- and post-intervention periods. This is most clearly explained graphically - Figure 3 below highlights the periods from which we had usable data before and after NLT installations (which occurred in the central white period). We only include households who have a good amount of winter usage captured between their 2 or more readings in both periods.

*Figure 3: Each horizontal line (one per household) represents the period over which we have data. Households are selected to ensure we have enough winter gas consumption data before and after NLT installations, which all took place in 2015*

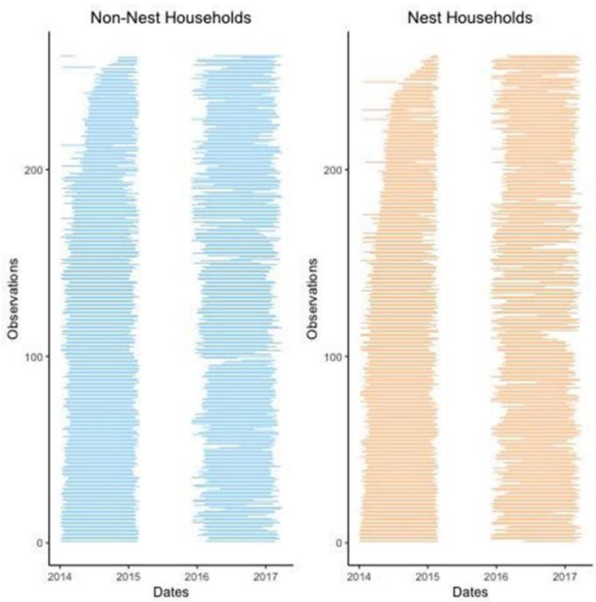


The downside of creating a more precisely defined panel dataset is that due to the very high number of estimated readings (which we discarded) and the strict conditions imposed, the sample size is greatly reduced. We consequently anticipated being underpowered, and therefore adopted an exploratory approach, running multiple analyses across multiple samples to look for trends and consistency in results rather than expecting to find significant findings. We assigned various criteria to the sample selection process described above, the least stringent being *<at least 2 readings and 20 days of winter data captured in both periods>* and the most stringent being *<at least 3 readings and 30 days of winter data captured between them in both periods>*. This created 6 different samples consisting of between 261 and 375 NLT households matched to either the same number of control customers (through 1-to-1 matching) or to a few thousand control customers (using 1-to-many matching). Given our original sample pool of ~4,200 NLT households and 27,400 control households, these small final samples reflect the very sparse and poor quality data.

For each of these 6 samples, we run 6 different regressions (36 specifications in total), each based on slightly different assumptions, and inclusion/exclusion of different covariates – these are detailed further in the full report and appendices.

**Study 3 findings**

- **Average savings in annual household gas consumption of 5.1% across 36 analyses of 6 samples** (p>0.05, 95% CI ±9.4%), compared to aggregate of other heating controls.
- **Very high levels of uncertainty, treated as exploratory research.**

With such sparse data and small samples the results are generally underpowered and therefore do not present statistically significant findings in the majority of tests. However, we find good consistency in the direction of effect, with 33 of the 36 models showing energy savings, averaging at 5.1% (±9.4%). Figure 4 below summarises the range of savings estimates across the 36 specifications, with the length of horizontal line indicating the 95% confidence intervals. Negative values indicate an energy saving, but only those whose confidence intervals do not cross the zero line are statistically significant – since we did not perform any corrections for multiple comparisons, these 3 significant results could well be arising by chance. We are therefore more interested in the general trend of these findings than in claiming statistical significance or the superiority of a particular specification or sampling strategy.

*Figure 4: Summary of savings estimates (with 95% confidence intervals) from Study 3. Note the very wide confidence intervals around each point estimate.*

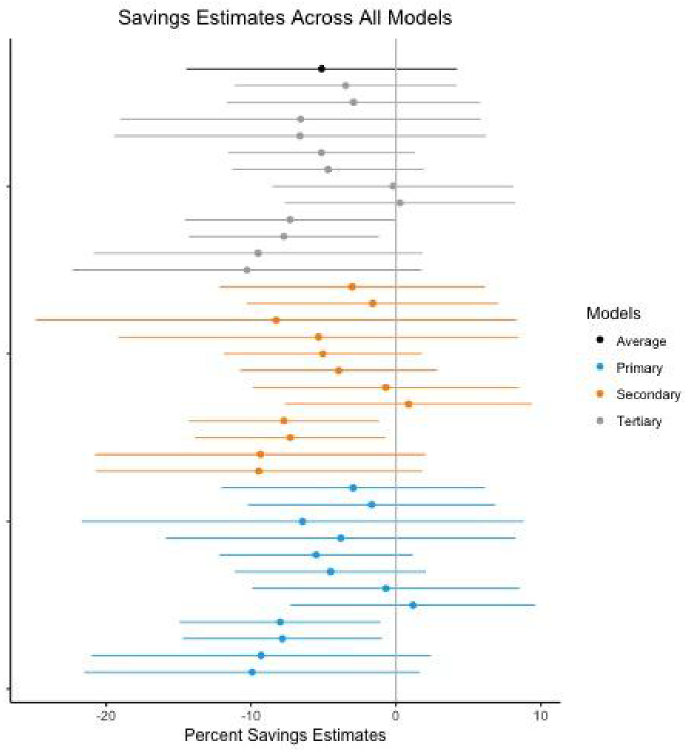


**Study 3 conclusions and caveats**

Though the results from study 3 are not strong enough on their own to confidently conclude energy savings, they are commensurate with Study 1 and Study 2 findings, and if we accept from those studies a prior assumption that the NLT does generate savings in the region of 5%, then Study 3 certainly provides no reason to doubt this and adds further supportive evidence.

The primary caveat is that these analyses are underpowered, and with extremely noisy and sparse data, the confidence intervals for each result (and the variation between results) are very large. Moreover, though the relative consistency with which these 36 results show an energy saving is encouraging, it is worth noting that these are not 36 separate replications, since they all draw from a similar sample pool (~60% of the NLT households are common to all 6 samples), and all model specifications have some similarities.

## Study 4 – Technical summary

### Study 4 methodology

We ran an independent evaluation of a quasi-experimental study run by Nest, in which 10,372 Npower households with NLTs were sent an invitation (via the thermostats and their smartphone accounts) to accept the Seasonal Savings feature. Seasonal Savings is an opt-in feature which aims to find incremental efficiencies in users' heating schedules in a manner which is likely to be acceptable to the user (for example, turning the temperature down fractionally overnight, and more aggressively than it ordinarily would when the home in unoccupied). This sample is compared to a matched cohort of non-Npower customers, also with NLTs, who are not invited to partake in the Seasonal Savings programme. This experimental design reflects the fact that Seasonal Savings is only offered through Nest's energy partners, including Npower, and it offers an opportunity to isolate the additive benefit of the Seasonal Savings algorithm beyond the standard NLT functionality.

Since our control and treatment groups both have NLTs, before and after the algorithm is offered and deployed, we are uniquely able to draw upon data directly from the thermostats themselves. As such the primary outcome measure is 'heating hours' – the number of hours per day the boiler is firing. This provides a very frequent, robust and direct measure of Seasonal Savings' impact on the heating system's energy consumption.

Nest ran the trial and undertook the matching of the two groups. We were provided with the raw data from the thermostats and ran an independent analysis, in addition to scrutinising the matching procedure. The intervention began on January 17 2017 and finished on May 8 2017, with pre-intervention data from November 1 2016 to January 17 2017.

We undertake two main analyses – first, an 'intention-to-treat' (ITT) analysis in which we measure the impact of the intervention across the whole eligible cohort sent the invitation. This average savings estimate therefore includes the impact on those who did not opt-in, i.e. we estimate the average effect of the offer, rather than of the algorithm itself. Given the nature of this matched, quasi-experimental design this is the most robust analysis, and is arguably most relevant to policy-makers wanting to know the average impact across a whole population in a world where the algorithm remains optional.

However we also want to identify the average savings on those who actually adopt it. There are two potential ways to do this: One approach is to analyse only those who opted in, compared to their matched control customers. However, this introduces self-selection bias in the treatment group, since acceptance of the algorithm is not random and matching was not done on propensity to opt-in. Instead, we therefore assume that all savings occur on those who opted in, i.e. no energy savings are associated with merely receiving the invitation for, but not adopting, Seasonal Savings. With this assumption accepted, we can generate an estimate for the average energy saving per opt-in by dividing the ITT estimate by the proportion opting in.

We find that around 10% of invitees were ineligible, principally due to having no live account and therefore being unable to accept the invitation. These are excluded from the analysis, on the basis that they could never have been treated. 84.6% of the remaining customers chose to opt in – notably high for an opt-in service.

We undertake balance checks to determine the success of the matching, and find a small but statistically significant difference in weather, with the Seasonal Savings group exposed to slightly warmer temperatures. This difference is small, with small differences able to reach statistical significance due to the large sample and frequent data, and thus the very high power of the study. The difference is also consistent before and after treatment. Controlling for this difference in the regression analysis therefore does not present a major issue. We find good balance on the main outcome measure (heating hours).

Our primary analysis provides an estimate for the reduction in heating hours. As a percentage reduction, this can also be interpreted as a saving in gas consumption of the heating system (since the heating system's gas consumption is a linear function of the boiler burn time, since the NLT does not materially influence the modulation of the boiler). We also convert these savings estimates into absolute savings in kWh, by first deriving an estimated kWh of gas used per heating hour by drawing upon smart meter gas data from Study 2.

The regression analysis uses a panel dataset with daily heating-hour data covering the pre- and post-intervention periods, controlling for thermostat-level weekly fixed effects to increase precision. Fixed effects are a flexible way of accounting for individual differences in preferences for heating as well as inertia in changing the set-point temperatures, and we find weekly fits better than monthly and is better than splitting the trial into pre- and post- periods only. We also control for weather (heating degree days), and account for whether the reading relates to a weekday or weekend. We remove the week over which the invitation was issued, as acceptance of the offer is spread over a number of days, with that week neither belonging in the pre- or post-intervention period.

**Study 4 findings**

- **Intention-to-treat estimate finds an average 3.8% saving** (p<0.01, 95% CI ± 1.0%), of heating-system gas use among whole target population, or **2.8% of total household gas consumption[30].**
- **This equates to an average of 4.5% saving** (p<0.01, 95% CI ±1.2%) **among those who opt in[31]**, or 3.3% of total household gas consumption[32], assuming no savings are achieved among those who didn't opt it.

In our ITT analysis we observe a significant reduction in heating hours per day (and thus gas consumption) of 3.8% (p<0.01, 95% CI ± 1.0%), in absolute terms equating to 0.151 hours per day of boiler burn time. As noted this ITT estimate represents the average savings across the full eligible group including the 15% who did not opt in.

We find an average saving of 4.5% (p<0.01, 95% CI ±1.2%) among those who opted in.[33] In absolute terms this is 0.178 heating hours per day. This estimate rests on the assumption that all observed savings occur among those who opted in, with no savings occurring in the 15% who received the invite but did not opt in. If this assumption is incorrect, any modest savings which do occur due to the invite would reduce this estimate (though only very slightly due to the high opt-in rate).

Each of these findings describes the savings from the heating system gas use. Expressed as a proportion of total household gas consumption[34], this equates to 3.3% among those who opt-in, and 2.8% across the whole targeted population.

---

[30] Assuming 74% of household gas consumption is for heating, as per Ofgem's data (72%) and our data from Study 2 (76%).

[31] 85% of eligible customers opt-in. Thus, 3.8%/0.85 = 4.5%.

[32] Assuming 74% of household gas consumption is for heating, as per Ofgem's data (72%) and our data from Study 2 (76%).

[33] Simply, 3.8% / 0.846, correct only if the assumption holds that no savings were achieved by the 15% sent the offer but not opting in.

[34] Assuming 74% of household gas consumption is for heating, as per Ofgem's data (72%) and our data from Study 2 (76%).

In monetary terms this equates to £18-26/year for typical medium-larger homes (among those who opt-in), in addition to the savings from the standard NLT features.[35]

We also draw on Study 2 smart meter data to estimate the average kWh of gas used per hour of burn time, finding it to be 16.1kWh per heating hour. This allows us to estimate the absolute savings as 2.24 kWh/day for the ITT estimate, or 2.66kWh/day among the opt-ins.

Figure 5 below illustrates the raw data, showing the reduction in heating hours after the invitation is sent.

*Figure 5: Heating hours (boiler burn time) per day for those who did (treatment) and did not (control) receive the invitation to Seasonal Savings*

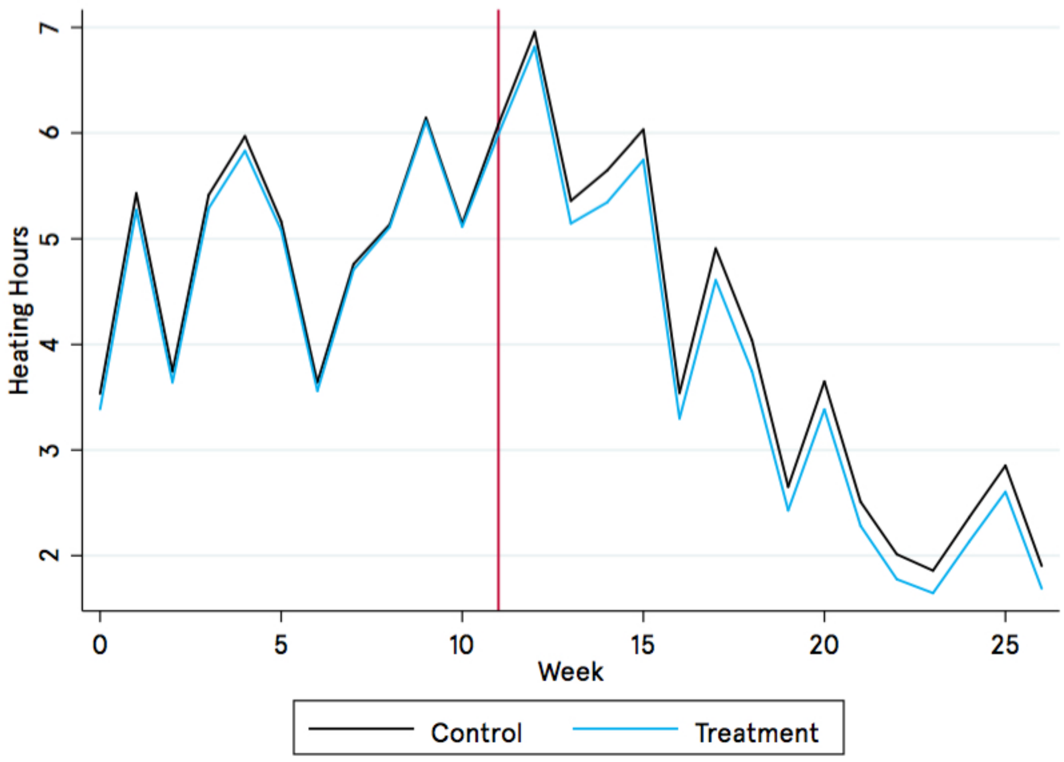


[35] As per the other cost estimates in this report, based on Ofgem's definition of a median home (12,500kWh/year), and larger home (75th percentile, 18,000kWh/year). Unit price of 4.35p/kWh.

**Study 4 conclusions and caveats**

Study 4 provides convincing evidence that the optional Seasonal Savings algorithm is able to find additional savings beyond those delivered by the basic functionality, as identified in Study 2. Despite the quasi-experimental nature of this research, thanks to the reliable and high-frequency data and large sample size, we are able to identify with reasonable precision the magnitude of these savings. The high opt-in rate is also noteworthy, and of relevance to policy makers who may doubt the benefit of an optional feature.