

Final report on the Amphiro-PWN-study

Effects of Real-Time Feedback on Hot Water Use



Date: March 31, 2016

Status: Draft V0.99

Authors: Thorsten Staake, University of Bamberg and ETH Zurich
Verena Tiefenbeck, ETH Zurich and Bonn University
Samuel Schöb, University of Bamberg
Anna Kupfer, University of Bamberg

Contact information

For research-related questions: thorsten.staake@uni-bamberg.de

For PWN-related questions: donald.mollee@pwn.nl



Bits to Energy Lab

Table of contents

Abstract	3
1. Motivation and study objectives	4
2. Introduction of the feedback device	7
3. Study design and execution	8
4. Results	11
5. Implications and conclusion	15

Abstract

Hot water consumption is the second largest position in a household's energy budget, only out-ranked by space heating. The average household uses more energy for water heating alone than for lighting, refrigerators, and consumer electronics combined. The nexus between hot water and energy use, however, is barely known among consumers, and consequently hot water is barely considered as a strategy to conserve energy even among very environmentally concerned citizens.

In a field study among 637 Dutch households, we investigated how real-time consumption feedback – the provision of information directly in the shower on the amount of water and energy used – influences hot water use. Participants obtained a smart shower display, which they were able to install by themselves between shower hose and showerhead. Data from 73'977 showers were recorded in a monitoring phase over three months starting end of August 2015.

In a baseline phase (feedback devices installed, consumption feedback deactivated), the average consumption was 3.2 kWh and 54 liters per shower. With feedback information (water and energy use plus energy efficiency rating), participants saved on average between 19% and 21% of their energy consumption in the shower. Absolute savings per shower amount to 0.6 kWh. The saving effects are statistically highly significant and stable over the entire intervention phase.

Projected to one year, a three-person household in the Netherlands (with 0.85 showers per person and day) saves on average 561 kWh of heat energy and 8.7 m³ liters of water. For the same household size, monetary savings amount to 86 EUR per year, leading to an amortization of the device within less than one year. From an investment perspective, abatement costs are below 0.05 EUR per kWh saved for an average 2.3-person household (for comparison: 0.10 EUR per kWh for large PV installations).

The study confirms the previous results of field trials conducted in Switzerland (together with Mobilier Insurance) with savings of on average 22-24% and 100 EUR respectively. In hotel, even when the user does not pay for water and energy, savings of on average 20% have been observed in trials conducted jointly with Mobilier. Heavy users and young individuals even save more on average.

Overall, the study confirms that personal real-time feedback on a specific, energy-intensive behavior on which the feedback receiver has a strong influence produces very large saving effects. Moreover, for the application in the shower, the measure has an excellent cost-benefit ratio and is applicable to the vast majority of households (home owners and tenants).

1. Motivation and study objectives

The water-energy nexus

In the Netherlands, water heating uses 2,460 kWh of energy per year and household (see Figure 1), exceeding the total amount of energy used for cooking, lighting, and ICT combined. While space heating, lighting, and electric appliance have received much attention in the past (e.g., in form of subsidized roof refurbishments, the prohibition of incandescent light bulbs, efficiency ratings, etc.), water heating only recently moved into the focus of attention among policy makers and firms as an important field for improvement.

Water heating deserves special attention in environmental campaigns, as residential water heating in Central Europe primarily relies on oil, gas, and electricity. This makes (hot) water generation very carbon intense, and measures to reduce the demand particularly worthwhile to pursue.

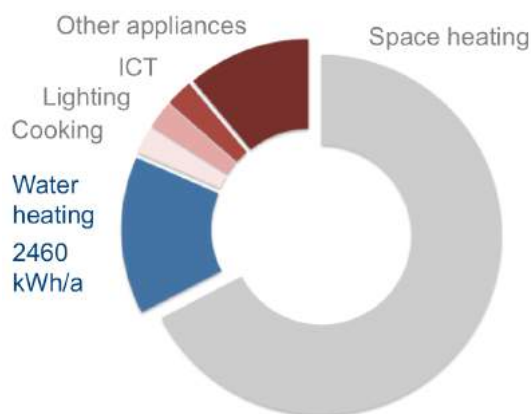


Figure 1: Consumption per household by end use¹

Despite its importance, the nexus between hot water and energy use is barely known among consumers. While dedicated accounts for space heating, smart electricity meters with in-home displays, enhanced electricity bills, etc. increasingly inform consumers and help them build up a general understanding on those domains, energy for personal hot water use is barely communicated to the end user. Consequently, hot water conservation is not considered as an energy conservation target even among very environmentally concerned citizens, leaving large saving po-

Water heating is the second largest position on a household's energy bill.

Water heating is very carbon intense.

The figure for water heating does not include hot water generated by dishwashers and washing machines (part of appliances) and water heated on the stove (part of cooking).

Citizens are barely aware of the large amount of energy that goes into water heating.

¹ Source: European Environment Agency (2012) / Split of ICT & Lighting inferred from bdew (2010).

tentials untapped. The understanding that drinking water is abundant reduces the intention to conserve hot water even further. Nevertheless, we expect a growing attention paid to hot water in the coming years: Its share among total residential energy demand will grow, as better building insulation and efficiency standards drive down energy consumption in other domains. In newly built, energy efficient buildings, water heating even exceeds the amount of energy used by space heating.

Shower water consumption as an energy saving target

Most of the hot water (about 70%) is consumed in the shower. An average shower requires 3.2 kWh (including losses) in only 8 minutes. There is no other activity in an apartment with a comparable power level (i.e., that uses more energy per time). At the same time, showering is an excellent target for conservation measures: First, the energy intensity allows for large absolute savings, and the short time interval of each shower only requires a short period of attention of a consumer to engage in energy conservation behavior. Second, showering is a daily routine of almost all citizens, so campaigns have the potential to address virtually all households. Third, only a very small fraction of the citizens knows how energy intense the daily shower is. Thus, information campaigns may leverage a steep learning curve. Forth, consumers have a huge influence on their shower water use. Cutting one's shower duration by only one minute leads to savings in the order of 13%. Fifths, shower consumption is something that can be easily attributed to one individual person (unlike electricity that is partly determined by other household members and by devices that have a base load), making it possible to provide persons-specific, very targeted and thus powerful interventions.

Supporting consumers with feedback interventions

Feedback interventions inform users about their behavior and/or its outcomes. The idea behind the concept is that the feedback receiver gets information on the outcome of a behavior she performed in a way that is easy to remember (often numerical or symbolic). A step counter is a prominent example for physical activity during the day. The feedback helps the receiver to evaluate her performance, compare it to previous days and to the performance of others, to set personal goals, to share the information among peers, etc. If provided in an easy to understand, motivating way, feedback interventions can have a large impact on behavior. Depending on the domain, feedback can support physical exercise, diet plans, the attempt to quit smoking, vocabulary learning,

Showering is extremely energy intense: 1 second in the shower requires as much energy as working on a laptop computer for 2 hours.

Showering is an excellent target for behavioral efficiency campaigns.

Feedback interventions inform the receiver about the outcome of a behavior.

teeth brushing – and energy conservation. In the energy domain, feedback interventions have been shown to lead to effects that are equivalent in effect size to large price increases.² The concept of feedback is often applied in order to curb electricity consumption (e.g., in the form of in-home displays with data provided by smart electricity meters or in the form of paper-based home energy reports) or as means to reduce fuel consumption (in the form of fuel gauges on a car's dashboard).

Feedback can be categorized depending on its timeliness (e.g., monthly vs. daily vs. directly after a behavior has been performed vs. already during a behavior). It can also be categorized according to its specificity (e.g., provided for a single behavior vs. for aggregated behaviors / provided per person vs. per household). In general, timely feedback (ideally given during the behavior) and person-specific feedback is more likely to yield large changes in behavior.

The major influence citizens can have on their shower consumption, the limited awareness among consumers (i.e., the steep learning curve), and the possibility to provide person-specific feedback during the behavior is performed makes showering an excellent candidate for feedback interventions.

Study objectives

The research goal of the study was to quantify and to better understand the effect of real-time feedback on shower behavior. We wanted to find out (1) how feedback changes the amount of hot water and thus the amount of energy consumed, (2) if the effects are stable over time, and (3) if specific subgroups of the study participants save more than others.

Additional goals from PWN have been to learn (4) how reliable the feedback devices worked in the field, (5) how the participating household responded to the devices, and (6) if the ratio between cost and saving per device is competitive.

The objectives have been addressed in a large-scale field study involving 637 Dutch households in 2015. The study was conducted by a research team located at the University of Bamberg, ETH Zurich, and the University of Bonn. PWN at Velsbroek financed the study and supported its implantation.

Real-time, person-specific feedback is expected to be especially powerful.

Showering is an ideal target for feedback interventions

The main study objective was to investigate the effect of real-time feedback on hot water use and the practicability of the feedback device.

² Source: Allcott, Hunt (2011). Social norms and energy conservation. *Journal of Public Economics* 95(9), 1082-1095.

2. Introduction of the feedback device

The central element of the study was a small measurement and display devices (“amphiro b1”) that provides real-time feedback on energy and water consumption in the shower (Figure 2). For the study, the 637 participating households installed study version of the device between their showerhead and shower hose. The devices turn on automatically as soon as water flows and provide information on water consumption (in liters) and energy consumption (in [k]Wh) for each individual shower together with an efficiency rating from A to G (A indicating a very energy efficient shower). The information was provided on a per shower basis (starting at zero for each shower extraction), and extractions that shortly followed each other were grouped into one larger extraction (to provide one number per shower even if the user makes short breaks, e.g., to shampoo her hair).

The smart shower meter “amphiro b1” provides real-time feedback on water and energy use directly in the shower.



Figure 2: Schematic representation of the feedback device amphiro b1

In greater detail, display area 1 toggles water temperature and the energy efficiency rating. When no water flows, the display remains on for about two minutes, waiting whether the shower is being continued and allowing the user to still see the final shower results. Display area 2 shows water consumption of the ongoing shower extraction during the shower, and toggles between water consumption and energy consumption when no water flows. Display area 3 provides an emotional representation of the efficiency rating (a polar bear standing on a melting ice floe; the size of the ice floe is coupled with the energy efficiency rating). Area 3 does not provide additional information, but has some emotional value and appears to make user talk more frequently about the device.

As the device automatically turns on whenever water flows, it does not require user action to start the measurement process. Nor does it require batteries, as it generates its energy from the water flow; this increases the environmental appeal of the device and makes battery changes abundant. The pressure drop is small – the device does not reduce the water flow in a very noticeable way. Unlike mechanical flow restrictors that face much criticism for restricting consumer choice and shower comfort in a paternalizing way, the feedback device lets users freely choose the flow rate and duration of their showers.

The device does not require a battery as it generates its energy from the water flow.

Besides displaying the information outlined above, the amphiro b1 stores (for each shower) the duration, volume, number of interruptions, and average temperature per extraction. Overall 249 such data sets fit in the memory. Base temperature (for the energy calculation) was set to 12°C. Device accuracy is +/-6% between devices at a flow rate of 12 l/min, with a much smaller deviation for repeated measures of one device. Flow rates can be between 2 l/min and 22 l/min.

The devices were built in Austria according to common environmental standards and use only drinking water compliant materials for all water-conducting parts.

The installation does not require any tools (see Figure 3). amphiro b1 is compatible to all customary 1/2" hand shower systems.




Figure 3: Installation of amphiro b1

3. Study design and execution

Experimental design

The study was organized as a field experiment in order to examine the effect of the feedback intervention in the real world (i.e., not in an artificial setting in a laboratory). Participants were randomly assigned to two

different groups³, the so-called treatment and the control group, which received group-specific devices. The devices handed out to the control group displayed only information on water temperature (i.e., no feedback on water or energy use). The devices given to the treatment group also displayed only water temperature during the first N*10 showers (referred to as baseline phase; N describes the number of household members using the shower), but thereafter automatically switched to feedback mode (the intervention phase). In the intervention phase, the devices provided the full set of real-time feedback on water and energy consumption as described in Section 2.

This design is referred to as randomized controlled trial with baseline phase. It allows us to investigate changes in consumption once the intervention of interest (here: feedback on consumption) becomes active by observing the difference between baseline and intervention phase. Moreover, by observing the control group, the study design also allows us to subtract non-intervention related influences (such as changes in outdoor temperature or changes in the behavior that stem from the feeling among the participants of being monitored in a study). The study design is illustrated in Figure 4. An online questionnaire was conducted both at the beginning and at the end of the study.

The field experiment was organized as a randomized controlled trial preceded by baseline measurements.

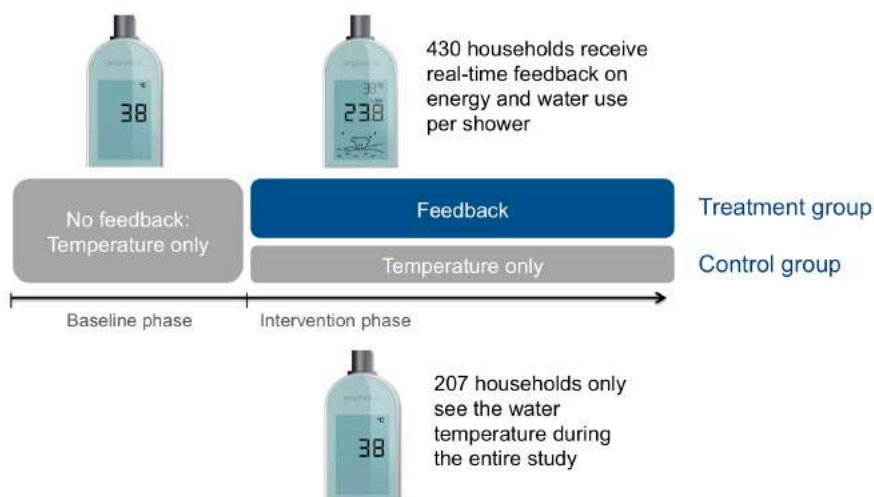


Figure 4: Experimental design

³ Randomization was performed while considering the target group size.

Recruitment of participating households and data collection

637 households participated the study. They were recruited among PWN employees including subsidiaries, PWN customer panels, PWN volunteers, and a group referred to as nudge panel that was available for research purposes. Prior to the first questionnaire, a short online survey was conducted to identify households that anticipated to relocate or to be absent for longer periods during the study or who had head showers (where the feedback device could not have been installed). This was done to avoid distributing devices to households that could not complete the study. Participation was voluntarily (“opt-in”) and free of cost to the participants.

637 households participated the study

Shower data was collected over a period of three months. Participants were asked to install a smartphone app that collected the shower data from the feedback devices and uploaded the retrieved data to a cloud server for subsequent analyses. These steps required a compatible smartphone with Bluetooth 4.0 connectivity. (iPhone > 4S and selected Android phones). In case of problems during data upload, the research team sent return envelopes and asked the participants to return the devices via mail. PWN employees also had the opportunity to drop the device off at the PWN headquarters. The research team then read out the devices, set them normal operation mode (so that control group participants received consumption feedback from then on) and returned the devices to the households. The process steps are shown in Figure 5.

Data was collected over a period of three months.

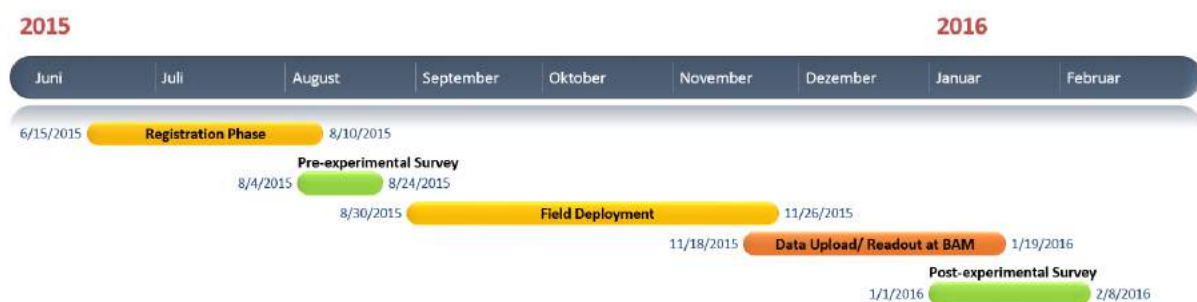


Figure 5: Study timeline

Return rate and data collected

Out of the 637 participating households, 503 provided data either by using the smartphone app or by shipping the device back for readout by the research team. The return rate of 80% can be considered as very good. In total, the datasets includes 73'977 shower extractions. From

In total, 63'206 data points were available for the subsequent analysis.

these, 63'206 extractions could be used in the subsequent analysis.⁴ This makes the dataset one of the largest ones covering shower behavior in the world.

4. Results

Graphical representation

The plot of the measured data nicely illustrates the effect of real-time feedback on energy consumption (Figure 6). We describe it step-wise to ease its interpretation.

The two lines shows the mean energy consumption per shower over the course of the study of the two groups (blue = control group, red = treatment group). During the **baseline phase** (no feedback, from 0% to 10% of study completion), the energy consumption of **control and treatment group participants** is almost identical, showing us that the random distribution of the participants into the two groups worked well. This is an important indicator that the participants are similar regarding important characteristics, and it increases the level of confidence that the effects observed in the subsequent intervention phase can be attributed to the intervention and not to group-specific differences. The first data point is noticeably lower than the rest. We assume that is because many participants who installed the feedback device and tried it out with a smaller water extraction, without actually taking a shower.

During the **intervention phase** (between 10% and 100% of the study completion), **control group participants** (blue dots) continue to see only water temperature. The consumption increases over time, as indicated by the upward slope of the blue line. We attribute this trend to the Hawthorn effect: at the beginning, participants “feel observed” and thus take shorter showers than they usually would; over time, they get used to the device and return to their normal shower habits. This is not interfering with the study results, as the effect is present for both groups.

With the onset of the **intervention phase** (feedback is shown for the first time, study completion rate 11%), **treatment group participants** (red dots) immediately reduce their energy consumption. This decrease is at-

Without feedback, both the control and the treatment group consume the same amount of energy.

The upward slope is attributed to the Hawthorne effect. It does not reduce the absolute size of the saving effects

The immediate effect of feedback is easily visible.

⁴ Most of the non-usable data points came from households where the number of showers exceeded the internal device memory, leading to corrupted data. For the sake of brevity, we do not discuss outlier removal here but include a discussion in an extended report to PWN. The detailed analysis shows that outlier removal does not affect control and treatment group differently.

tributed to the feedback intervention. Savings are represented by the difference between the two trend lines. The trend lines are almost parallel: the treatment effect remains constant during the experiment. If there is a change, then the gap seems to widen; that would mean that the savings even increase the longer the participants receive feedback.

The following subsection quantifies the effect size. The graphical representation, however, already conveys the large effects size from real-time feedback on shower behavior.

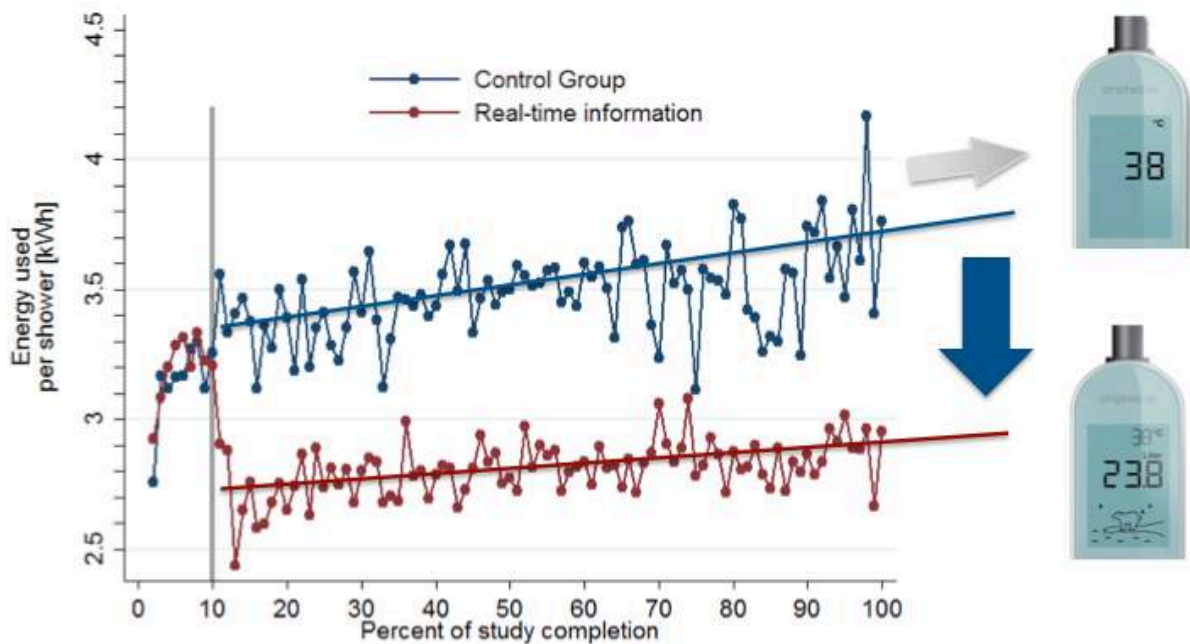


Figure 6: Feedback effects on per-shower energy use

Calculating the effect size with a difference-in-differences model

In order to quantify the effect size, we calculated the changes in consumption with a difference-in-differences (DID) analysis. A DiD analysis compares the mean energy use of the two groups (control and treatment) during baseline and during the intervention phase. This relatively simple approach has the advantage over more sophisticated regression models that it is more straightforward to understand and verify.

For a DID analysis, one derives the difference between control and treatment group during the baseline phase and subtract from it the difference between control and treatment group during the intervention phase. In our case, this reveals average savings per shower of 0.64 kWh, or 20.8%. The analysis is illustrated in Figure 6.

The difference-in-differences model shows savings per shower of 0.64 kWh.

Considering also the results of a regression analyses, we regard per shower savings of 0.6 kWh as very reliable estimate.

An alternative to DID analysis is to estimate a more complex regression model. Using a fixed effect regression model, we found the savings to be 0.55 kWh, or 19.6%. Given the inherent error margins of field studies, this virtually the same result as shown by the DID analysis.

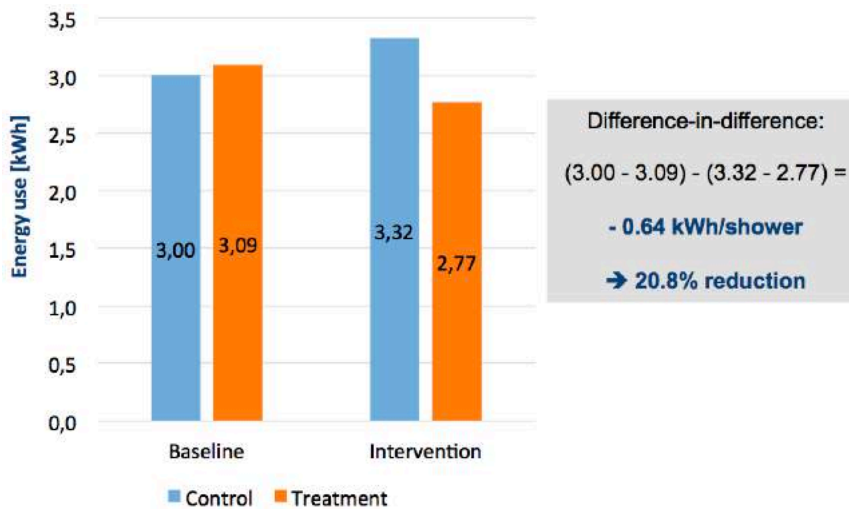


Figure 6: Calculation of the effect size with a difference-in-differences approach (energy use per shower, no minimum threshold filter)

The energy savings almost completely result from a reduction in shower duration. The treatment group only slightly reduced the flow rate, and took their showers at almost the same temperature. This is not very surprising: reducing the duration of a shower by a minute or two is hardly noticeable given a human's sense for time at these scales, while a reduction of the water temperature would result in a severe loss of comfort. In addition to the energy savings, the average amount of water saved per shower is 9.3 liters. While the energy savings account for a considerable share of the households total energy budget, the water savings only account for a very small proportion of total water use.

Stability of the effects

For both, the difference-in-differences and the regression model, the results are robust to different filters (e.g., excluding or not excluding extraction with unbearably hot shower water temperatures or very short (4.5 liters) extractions that probably result from bathroom cleaning rather than from showering). Moreover, the regression shows no significant time trend in the effect size, indicating that the savings remain constant.

In addition, households save on average 9.3 liters per shower.

The saving effects are statistically highly significant and stable over time.

Attempt to explain the large effects

Savings from smart electricity meters are typically in the range of 2% to 5% (Schleich et al. 2013, McKerracher & Torriti 2013), amounting to about 80 to 240 kWh per year. This raises the question why the saving effects observed in this study are much bigger. Several factors *may* explain the larger effects observed here: First, the feedback on hot water use is given already during the behavior is performed, allowing the user to immediately respond to it. Second, the feedback is person-specific and directed to a specific behavior, making it easy to understand and to relate to. Third, the feedback receiver has a high level of control: Unlike feedback on electricity, the energy consumed in a shower is not influenced other household members, base load from appliances, etc. and the valve controlling the consumption is within reach. And forth, the curtailment behavior requires only short time span of attention, as a lot of energy is consumed in only a few minutes. These factors appear to make showering an excellent target for feedback interventions.

The savings effects a much higher than for smart metering for electricity. It is easier for consumers to respond to feedback provided in the shower than to keep track of their electricity use that is influenced by many factors.

Total energy and water savings / abatement costs

So far, we have reported the savings per shower. The data suggest that a person takes on average 0.85 showers per day (or almost six per week). For the average Dutch household (2.3 persons), this results in savings of 428 kWh plus 6.7 m³ of water and wastewater per year.

Table 1 summarizes the energy and water savings and the abatement costs (investment per kWh saved) over a three-year period. Three years have been chosen as a conservative estimate for the device lifetime. For the abatement cost calculation, device costs of 60 EUR have been used. Note that the numbers only reflect the pure cost side: They do not include the households' savings on their energy and water bill. If those savings are taken into account, the costs per kWh saved are negative; the cost-benefit analysis is presented on the next page.

Household size	1 person	2 persons	2.3 person	3 persons	4 persons
Energy savings	558 kWh	1'117 kWh	1'284 kWh	1'675 kWh	2'233 kWh
Water savings	8.7 m ³	17.5 m ³	20.1 m ³	26.2 m ³	35 m ³
Investment per kWh saved	0.107 EUR	0.054 EUR	0.047 EUR	0.036 EUR	0.027 EUR

Table 1: Water and energy savings for different household sizes and abatement cost within three years

Monetary savings for different household sizes and heating scenarios

For individual households, the monetary savings can be estimated using current water and energy prices. For the calculation, we differentiate between two heating types (gas and electricity), which are both common and which have different cost of energy. Table 2 summarizes the monetary savings over a three-year period. Cells highlighted in bold indicate a return on investment in less than one year.

The following utility cost have been used for the calculation:

- Cost of drinking water: 1.83 EUR / m³
- Cost of waste water: no variable cost
- Energy cost gas: 0.079 EUR / kWh
- Energy cost electricity: 0.23 EUR / kWh

Household size	1 person	2 persons	2.3 person	3 persons	4 persons
Heating: Gas ⁵	60 EUR	121 EUR	139 EUR	181 EUR	241 EUR
Heating: Electricity	145 EUR	290 EUR	333 EUR	435 EUR	580 EUR

Table 2: Monetary savings for different household sizes within three years

5. Implications and conclusion

The results show that real-time feedback on shower water use leads to very large saving effects. This especially holds for households with more than one person. For families, the payback periods are less than one year, even at today's low energy prices for fossil fuels. The effects are also noteworthy given that the technology is enabling rather than patronizing by nature: It gives the users the opportunity to act in line with their intentions rather than relying on higher energy prices or restrictive measures such as flow restrictors.

Another strength of the feedback intervention is its compatibility with the vast majority of showers. Unlike most other technologies for heat energy conservation that suffer from split incentives (where home owners have to invest while tenants benefit from lower energy bills), its in-

Large-scale savings can be achieved with abatement cost of less than 0.05 EUR per kWh.

⁵ Note: The calculation uses average values for heating efficiency. Most likely, gas heating has a lower-than-average efficiency (savings are slightly higher than in the table), and electricity heating has a higher-than-average efficiency (savings are a bit lower).

stallation makes sense both for tenants and homeowners. This makes the technology well suited for large scale campaigns.

From the perspective of administrators of an energy efficiency campaign, the abatement costs (i.e., the investment necessary to save one kWh) are very low compared to other measures. When specifically targeting families (three-person households), the investment per kWh saved is 0.036 EUR.

A challenge related to the adoption of the feedback intervention remains. Consumers rarely relate hot water to energy use. Consequently, hot water conservation is not considered a saving target even among very environmentally concerned citizens. This makes marketing (hot) water saving technologies a challenge, probably requiring an enduring approach to raise the awareness among citizens.