# Flipping the Energy Data Iceberg

## Views into the Consumer Benefits and Energy System of the Future through Intelligent and Adaptive High Resolution Metering

**Authored by Nick Lange and Dan Fredman, VEIC; Mike Phillips, Ghinwa; George Zavaliagkos, Sense Labs**

**August, 2018**

# Table of Contents

# Abstract

If Automated Metering Infrastructure (AMI) data-driven savings today is what we can see above the water, then what's the 90% of energy savings iceberg hidden below? New answers are suggested by recent findings from approaches that bring device-level disaggregation, and smartphone ecosystem integration towards a more perfect union of technologies and a focus on the people they serve.

We assess next generation home energy monitoring systems—those that are using machine learning and intensive computation—in their effort to make the kernel of truth in the one-liner for home energy monitoring: "mean time to kitchen drawer" a thing of the past. We'll share a framework for energy stakeholders (utilities, policy-makers, etc.) to better understand and consider how this enabling suite of capabilities (high sampling rates, real-time low-latency feedback, and a consumer-friendly orientations) may affect their approach to accelerating market transformation through future research, pilot programs, services, and regulations.

The paper's perspective is anchored by a review of first-hand experiences and novel findings and results from a selection of case studies. It will include qualitative customer insights and large-scale energy analyses based on thousands of devices in the field. Although even the most aggressive forecasts indicate that we are likely a decade or two away from when most things in most homes are smart and connected, this review of the latest behind-the-meter products and services may help us start to better conceptualize the energy systems of the future that are driven by a new kind of EM&V: Engagement, Meaning, and Value.

# Introduction

The lighting cliff, duck curve, and utility death spiral have more in common than their flamboyant flair for memorable rhetoric. Each is a difficult energy efficiency problem whose solutions may best be found not from within the details of their specifics, but in following the counterintuitive strategy of "making the problem bigger."[1] Expanding the scope of concern can often bring into view new assets and resources that can be reorganized into tackling a larger systemic concern. We believe that this is the case in the hard and enduring problem of realizing significant societal and energy value by successfully engaging people with energy data.

[1] Popularized by General and United States President Dwight D. Eisenhower. Cited and advanced by the diverse interests of management consultant Peter Drucker, politician Newt Gingrich and scientist Amory Lovins.

Through early deployments of new energy disaggregation technology, supplemented by utility behavioral efficiency program integrations, we are beginning to glimpse the outlines of solutions to a surprisingly broad spectrum otherwise intractable energy systems issues. These promising bright spots of customer-centered, data-rich approaches offer a sharp contrast to the struggles of more conventional energy and efficiency economic levers straining and to serve diverse and dynamic market needs of the present and near future.

Not just victims of their own past successes, institutions which thrived by harvesting bumper crops of low-hanging technology and policy fruits (e.g. in lighting and utility rate decoupling) must now navigate increasingly complex savings opportunities and rapidly accelerating rates of change. New interventions, methods of measurement, resource management and regulatory structures may be required. All of which must benefit from economic efficiencies that can scale and advance at the speed of modern markets. What is an energy stakeholder to do?

A surprising and compelling answer is suggested by growing the problems of energy disaggregation far beyond the basics of accurate and reliable billing analyses. It now appears viable that heretofore unavailable relationships between customers and their energy data may soon be unlocked by innovations that can sow seeds in the fertile fields that fit into—even facilitate—the everyday interests of people.

# Differentiated Disaggregation

The topic of load disaggregation and how appliance level information can be used for driving efficiency gains has been explored in some detail (M. Zeifman, K. Roth, 2011) and (Armel, C. and Gupta, S. 2012), but we are not yet aware of studies of large-scale deployments of systems with detailed appliance-level breakdown of energy and the results obtained. Much of the work that has been done has focused on the possible benefits of a detailed energy breakdown and likelihood this could be used to drive positive consumer behavior. But, we believe that to have maximum benefit, and to sustain those benefits over time, it is crucial to have consumers who are actively engaged in how energy is being used in their homes.

We don't think this can be done based purely on an energy usage view—especially a historical energy view. Consumers just don't care enough about energy to spend much time thinking about it on an ongoing basis. Even if a consumer does care about energy consumption (based either on a cost-saving motivation or based on an environmental concern) and is willing to spend some time on it, this tends to be a short-term focus followed by a drop off in interest and attention. The reason is pretty clear – once you understand how energy is being used in your home, there is not a lot of reason to come back and keep checking. Ongoing awareness is important to handle degradation in appliance performance, or reversion to less efficient behaviors, but it is a hard sell to a person with many rational, and emotional competitions for their attentions. The initial period of understanding energy use in your home can be interesting, but the ongoing monitoring is not.

## Consumer Engagement

### Customer Engagement – Definition

Developing mutual understanding with customers over time. We offer this description to help clarify the difference in practice and experience from other more transactional terms like conversion. This reinforces the idea of forming a relationship that fosters trust which can play a role as catalyst for reducing costs and increasing benefits for all parties.

Moreover, the ability to confidently and specifically correlate these ongoing interactions with measured energy impacts can be critical to meeting evaluable standards of attribution that are required for most utility efficiency programs.

### Customer Engagement – Enabling Technology

So, we need another reason for consumers to engage with energy-related applications and we believe the key to this is to expand the notion of home awareness beyond just energy, and to give consumers a view into what is going on in their homes and their lives around them. This is a new orientation from the traditional utility energy data initiatives. This human-centered approach is a bit like a new glasses prescription in designer frames: it sharpens acuity, even if the new look takes some getting used to.

Though this perspective is new for many utilities, the general category of tech is not. This class of product fits well-enough in the existing taxonomical grouping in the literature for energy feedback devices (Bertoldi, Serrenho, Zangheri 2016) as the most "Direct feedback with 'Connected Devices' and Automation" seems suitable enough. Rather than mining utility billing data, these offerings largely self-serve from their own data sources: additional hardware is installed in the electrical panel of a home so that much finer data can be collected and computed on instantaneously for disaggregation and communication purposes. Supported by a smart phone app and web ecosystem of cloud-services, at present this kind of virtual sub-metering costs more than running analyses on existing AMI data, but costs less than, and is more flexible and resilient to change than circuit-level monitoring.

Sense[2] is an example of an application which does this.[3] Sense's home screen has both a real-time view of energy consumption (and production if a homeowner has solar) what devices are on and off, and also provides a scrollable timeline of device activity in the home. An example of the interface can be seen in Figure 1.

---

[2] Note that Sense is a product of Sense Labs, co-founded by one of the authors of this paper

[3] Other examples with a similar customer-facing orientation include *Smappee*, *Neurio* and others. For a more complete list, with annotations, see p. 7 of (Kelly 2017)

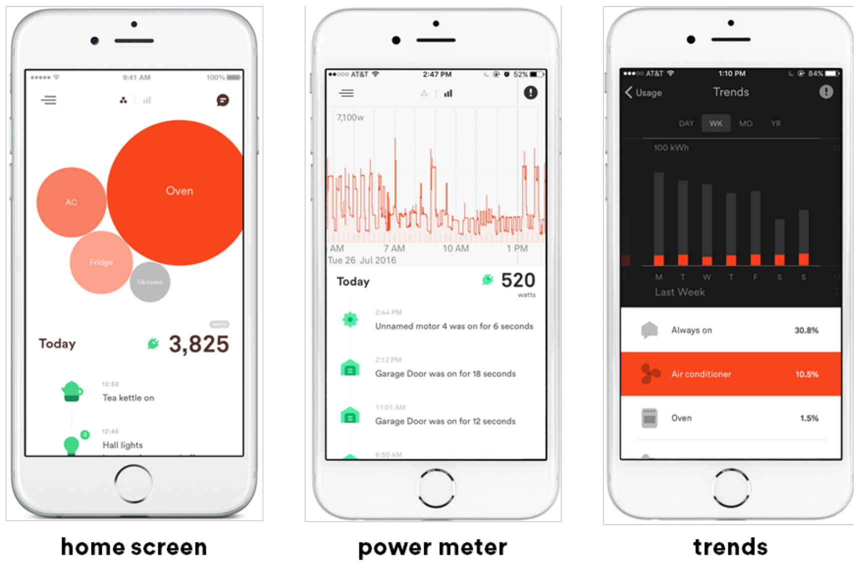


Figure 1: Screenshots of Sense application

In addition to the home screen, there are screens with a detailed real-time power meter, screens with various trends and statistics, and screens with device details.

While the application provides quite a bit of energy related information, it is the scrollable timeline of device activity which is the main driver of sustained engagement. Using this, homeowners are able to answer questions like “did I leave the oven on”, “is the sump pump running”, “did my kids get home from school”, etc. The use cases vary widely by homeowner, but for many homeowners, there is at least one thing that they want to know on an ongoing basis which they can learn from the application.

It is well known that most energy related applications—like most consumer mobile applications in general—start with engaged consumers, but then drop off to very small numbers over time.[4] While it is difficult to measure user engagement, we look at a number of metrics of application usage (monthly actives, daily actives, etc.) over time. Figure 2 shows the engagement numbers seen with Sense over time. Percentage weekly actives means the percentage of all users who opened the mobile application at least once in a seven day period. Weekly usage is expressed as the average number of sessions per week for those active users. The average session length was 51 seconds.

[4] Across “all industries” the average mobile app retention rate was 20% after 90 days. (Localytics 2017)

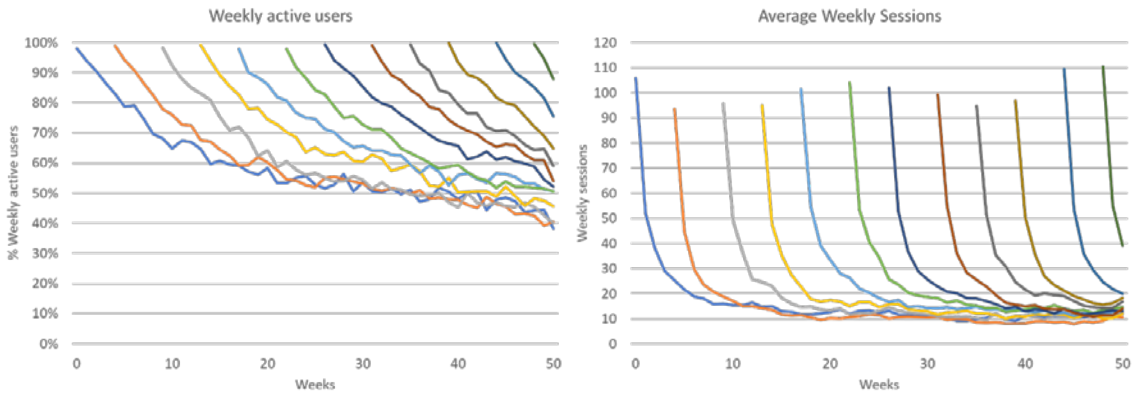


Figure 2: weekly actives and usage over time for monthly user cohorts for Sense's US install base.

These usage numbers were computed for cohorts of users based on month of install. So, the leftmost line is for all of the users who installed Sense in February 2017, the next line for all users who installed in March 2017, etc. You can see that there is the normal drop-off in usage over the first few months, but then usage mostly levels out to something around 40-50% of users using the application on a given week and those users are opening the application a bit more than 10 times per week on average.

How does this compare to typical levels? Unfortunately, it's difficult to say due to the sparse availability of quality energy engagement data. Whether as a result of the non-digital nature of the program (HERs) or for business purposes, most energy feedback applications haven't made this level of information generally available. One point of context, albeit one of only moderate usefulness, is in relation to the nine minutes per year (Accenture 2012) that the average customer reports engaging with their utility. Assuming an average app session lasts just under one minute, then the time Sense's customers engage in represents almost a 50x increase. For better or for worse, these numbers can't muster a score on the infamous disengagement "mean time to kitchen drawer" metric. Even after a year since installation, Sense is keeping a safe and steady distance from the 'drawer' so far.

fields that fit into—even facilitate—the everyday interests of people.

# Energy Details

We think the path to large and sustained energy efficiency gains is to combine the consumer engagement we are now seeing with specific and actionable insights and suggestions around how energy is being used in a home and how that usage can be improved. These intuitive and accessible hypotheses are suggested by multiple aspects of theory, supported by recommended behavioral science practice, and affirmed by early experience.

## Behavioral Savings – An Expanded Definition & Opportunity

For the purposes of discussion, we propose a definition of behavioral savings for energy efficiency programs as the energy reductions attributed to information-based programs and primarily measured by energy usage. Within this broad frame—agnostic on the nature of informational program, behaviors changed, and the particular evaluation methods—the amount of behavioral savings available expands to all energy reduction opportunities[5] that can be evaluably attributed to program-driven intervention.

That, though, is the rub: reliable methods for predicting, tracking and attributing energy impacts. The sheer breadth, and relative paucity of information about behavioral savings stymies conventional, measure-based characterizations and impact evaluations. Even the gold standard methods of most HERs programs are frustrated to satisfactorily explain the underlying energy end-use source of savings, expected lifetime, and future potential.

One particularly promising area of appeal for engaging apps and device level disaggregation approaches like Sense, is the new availability of entirely new dimensions of energy AND engagement information. Combined, these channels may provide paths not only to better understand the underlying appliance level sources of behavioral savings programs, but also their correlative causes required for program measurement and verification. Large scale experimental design, sampling, and surveys all have important roles to play, but the sum of their parts may be poised for rapid growth in understanding and capability.

## Behavioral Savings – A More Reliable Messenger

The vast majority of behavioral science tools rely upon the timely and specific conveyance of a tuned communication to the party of interest. (Ashby 2017) Thus, the successful outcomes depend to a rather outsized degree, upon both the messenger and the medium. For all the apparent power of injunctive social normative comparisons, if the source is corrupt and unreliable, it's hardly persuasive. And, if you don't even get the message in the first place, it hardly matters what it says!

In either case, there is a certain resilient strength in a sustainably engaging smartphone application like Sense's. The hope is that once consumers are using an application like this on an ongoing basis and trusting it to tell them things like when the garage door opened, or the dryer's done, that they will also trust it when it lets them know their air conditioner is not working correctly or that they should replace an old refrigerator. With multiple sessions steadily available each week, there are ample opportunities to appropriately try and deliver the message. Early indications are encouraging.

[5] The enlarged potential for savings realization through diff-in-diff RCT Home Energy Report programs is a significant, albeit somewhat underappreciated, contributor to the scale and staying power of HERs savings

# Pilot Program – Advanced Residential Intelligent Efficiency Services (ARIES)

In the fall of 2017, Efficiency Vermont[6] launched a pilot program to explore the viability of efficiency program portfolio savings of products like Sense. A detailed description of this initiative[7] is best left to future papers, but two early examples help to show the benefits of these levels of enhanced customer engagement, detailed data, and opportunities for analysis. Figure 3, below, shows the number of by program participant completions over time in response to two e-mails, an initial 'newsletter' style piece with multiple content modules, and a survey-specific reminder targeted to non-responders. Despite the lack of any incentive offer or deadline, the completion rate exceeded expectations, reaching 68% within two weeks. This rate rises to 94% when compared to the number of program participant opens for the initial 'newsletter' e-mail. As the pilot calls for multiple surveys over time, this level of response is a key indicator of success.

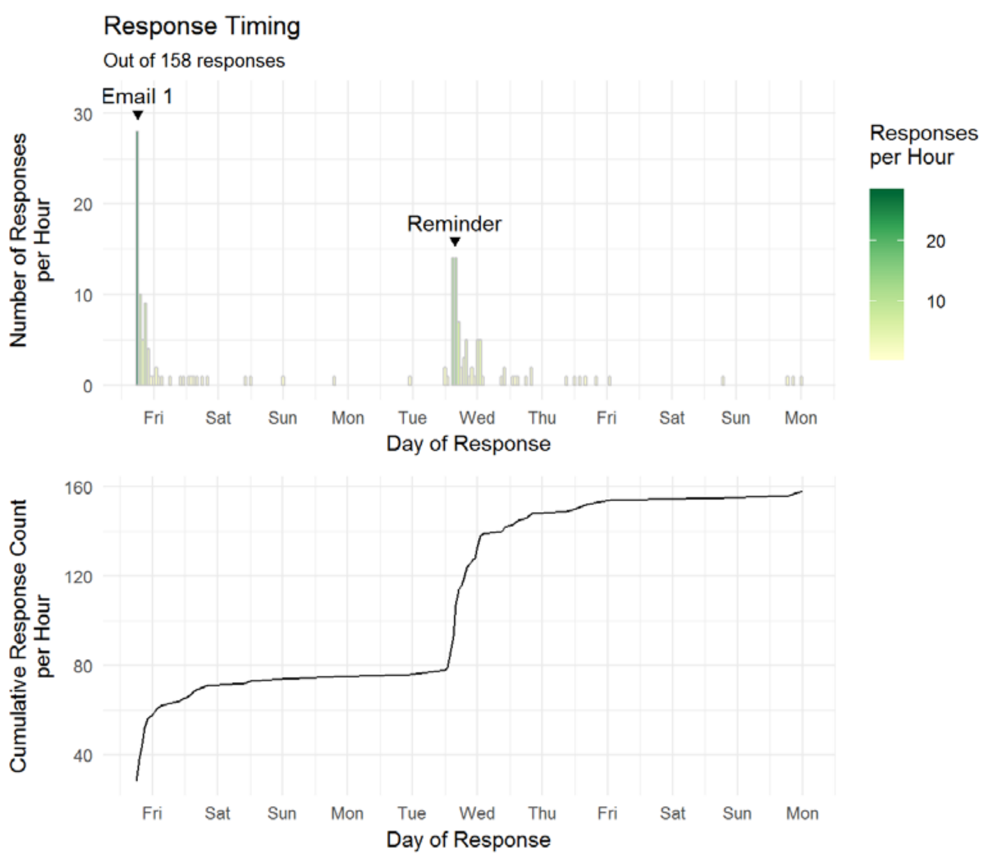


Figure 3: Survey completions over time. Note the clear visual evidence for the effect of the reminder.

---

[6] Efficiency Vermont is administered by Vermont Energy Investment Corporation (VEIC), an independent nonprofit energy services organization under an appointment issued by the Vermont Public Service Board.
[7] The pilot program is a highly collaborative effort, including multiple innovative integrations between the utility efficiency program, the product vendor, and the customers they serve.

High survey participation rates provide a more complete foundation for program delivery, customer service and satisfaction, but perhaps most critical are their contribution measurement and evaluation. Figure 4, below, shows two visualizations from early analysis of open-ended text responses to the question “Since you got your Sense, what small energy-related changes have you made?” In the upper graphic, the presence[8] of text associated with each keyword family ‘action’ is shown as a column for the 90 respondents self-reporting changes. The lower graphic provides an ordered frequency distribution for these actions and their combined presence.

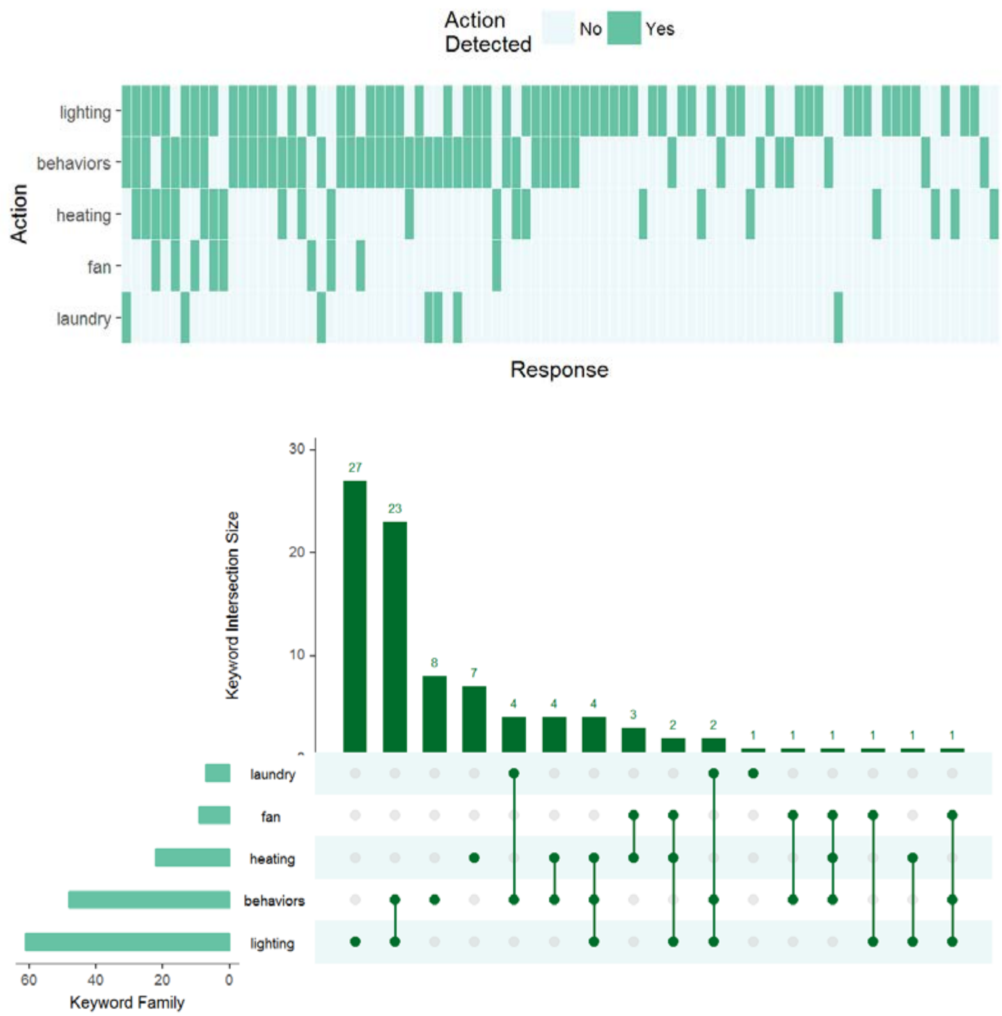


Figure 4: Concurrent detection (above) and sorted counts of groups of self-reported energy-related changes

[8] Free-text response fields suffer from negative cognitive effects like recency and availability biases but avoids a number of problems that arise from asking participants to select from a list of specific and pre-defined behaviors. As a result, the text must be iteratively examined for categorization and comprehensive inclusion through encoding keyword families and validation to catch misspellings, programming errors, and omissions.

Outputs from preliminary steps like these provide the basis for: segmented communications and personalized messaging, data model development and training, and improve program capabilities to identify, correlate, and appropriately attribute energy impacts of information-based programs.

## Behavioral Savings – Targeting Potential

The overall notion is that by using detailed energy breakdowns and comparing with other participants (taking into account location, house size, weather, etc.), that we can identify specific improvements which could be made in a specific home and then through email, in-app alerts, etc., encourage the user to make the change. Since we continue to get ongoing information about the detailed energy use in the home we can also verify on a house-by-house basis whether the change was made. Ultimately with a large enough sample size we will be able to do very fine-grained A/B testing of outreach approaches to determine most effective messaging - all based on this detailed knowledge.

While there are many possible detailed loads to go after, so far, we have focused on three specific cases: refrigerators, incandescent lighting, and "always on"[9] loads. These are things that Sense can detect with reasonable reliability and are well known targets for energy efficiency gains. For refrigerators[10] and incandescent lighting, the issue and remedy can be fairly clear to the end user. And, the economics of replacement are well known.

For "Always on", we track this by estimating the minimum load over a rolling 24-hour period. This minimum is robust to small dropouts or other noise effects and is meant to determine the baseline load for the home. In most cases this does match the things which are truly drawing constant power (routers, DVRs, radon fans, etc) but in some electrically busy houses may include some amount of variable loads also. Compared to refrigerators and incandescent lighting, this can be a trickier category to explain to users and to help them track down, so one of the efficiency challenges is how to help users with this.

Because we now have detailed data across homes, we are able to determine potential savings for each of these categories by looking at distributions of usage across homes. For each category, we are computing a metric of potential savings, by computing the savings if X% of the worst performing devices could be replaced with ones with the median performance of all the devices in the group.

This of course is not a perfect measure – we know there will be no realistic way to get all of these users to take this action, and if they do take an action, they may be able to replace with a new device which is better than the median, or get rid of the load all together (many of the inefficient refrigerators are older models kept around as a spare

[9] "Always on" is comparable to "phantom loads" and "home idle load." (Delforge, P., Schmidt, L, and Schmidt, S 2015) explores the topic in some detail, and estimates a $165 cost to an average US home.
[10] Sense's refrigerator identifications include a small portion of freezers – they are currently unable to distinguish them from refrigerators based on energy signals alone.

when the homeowner upgrades to a newer model – in some cases they may opt to just get rid of the spare rather than replace).

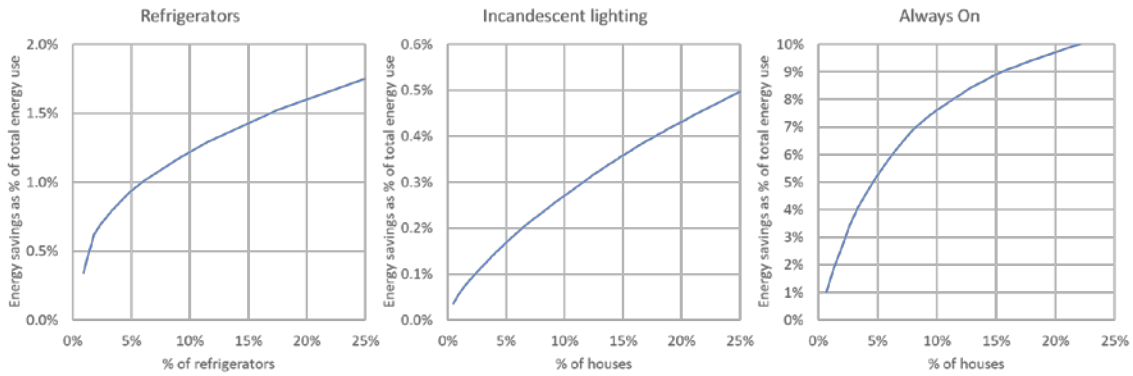


Figure 5: possible savings by category. For each, possible savings as percent of total power use as worse performing X% of devices are replaced with devices with median performance. VT pilot participants.

But, this does allow us to compare possible gains across categories. Figure 5 shows the potential gains expressed as % of total energy use (across the entire population) if the worst performing X% of devices are replaced with a new version with median performance (in the case of lighting, we assume incandescent bulbs are replaced with LED bulbs with similar light output).

From figure 5, it is clear that such well-established efficiency targets (refrigerators and incandescent lighting) seem to have largely run their course - at least for our pilot participants. Perhaps unsurprisingly given the population of interest—opt-in residents in a state with a long history of efficiency program success, there does not seem to be a lot left to be gained from these. On the other hand, always-on devices have been proliferating and are a clear target for energy reductions. Moving just 10% of the worst performing houses to match the median performance would result in an overall total energy savings of almost 8% of total usage. In addition, moving the median always-on usage down by 10% relative, would result in an additional 2.7% savings of total power.

## More Devices

While we have started with just three well-known categories (refrigerators, incandescent lighting, and always-on), we will also extend to a broader range of specific efficiency gains based on a broader array of devices. The biggest clear category is HVAC performance. This requires modeling not only electrical use, but information about location, weather, house size, and HVAC equipment type. This work is ongoing and we will apply the same user interaction and active measures as the modeling is more complete. We also believe there is a long tail of specific cases (dehumidifiers left on, roof heating coils left on, use of resistive heating[11], etc.) each of which are small, but when combined are another major portion of potential residential

[11] Beyond conventional electric heat (baseboard and spot heating units) this also includes less blatant sources embedded within water heaters and heat pumps.

energy savings. Because this is a diverse set, the detailed per device performance of a load disaggregation like what is found through Sense is needed.

## Creating (and measuring) Change

We are first testing just the change that is caused by users having real-time access to power usage along with the detailed appliance-level breakdown. We believe, and there is some reason to, (Kelly and Knottenbelt 2016) that even without active measures, there will be some gains and will use this as a baseline for the active measures.

As a next step in this collaboration, we will be testing and observing the effect of a variety of specific measures to reduce consumption in these target cases. Even through refrigerators and incandescent lighting only have modest potential average gains, we will include these also as clear test cases to build the customer relationship to this kind of messaging. The active measures will include:

- General emails to the users with suggestions and tips for how to save energy using the real-time power views and appliance level breakdowns
- Specific emails showing the user their energy use versus the population
- For users who have clear targets for reductions (a poorly performing refrigerator, or extensive incandescent lighting use for example), the email will also include specific suggestions for how improve and the economics of doing so
- New application features to allow users to set goals and budgets and track performance against these goals (including things like an alert if always-on exceeds some threshold over some period of time)

For each, we expect to be able to track the effect and iterate as needed. Since we are able to track specific consumption in specific homes, we can monitor for changes after employing each of these active measures above and iterate as needed to find and focus on the measures which work.

Most programs would like to follow the best-practice recommendations for the evaluation of behavior savings programs (Dougherty and Provencher 2013) which emphasize the need to focus on the specific behaviors of interest. But you can’t be very specific when only working with coarse energy data and little meaningful data on customer actions. In contrast, this new approach has far more to work with to better serve prevailing M&V guidance. So, the pilot aims to apply lessons from behavior science[12] not just to the program’s design, but also to the measurements of its performance.

This requires careful analysis to integrate and align: 1) Content of engagement triggers from in-app and e-mail messaging; with 2) Survey data insights into customer

[12] Although a large number of behavioral science insights are involved, the specific language (e.g. trigger, motivation, ability) and structural mechanisms of the Fogg Behavior Model (Fogg 2009) are used here to help more clearly convey the general idea.

motivation and capability; so that 3) Results from device level energy impact models are understood and properly accounted by energy program stakeholders.  It won't be easy, but it's now possible to do, and thus far more likely to succeed.

# Looking to the Future

The lack of evidence for significant energy efficiency benefits of 'standard' smart meter data-based disaggregation compared to aggregated energy feedback is as discouraging as it is inconclusive. Most of the work around load disaggregation has been focused on the possible benefits of a detailed energy breakdown and likelihood this could be used to drive positive consumer behavior. But, we believe that to have maximum benefit, and to sustain those benefits over time, it is crucial to have consumers who are actively engaged in how energy is being used in their homes. The approach being piloted in Vermont—an efficiency program deeply integrated with a product like Sense—brings such a massive increase to the quality and character of relationship to customers, and energy data; that it is hard to ignore.

The detailed and real-time information provided by an application like Sense is what we think is needed for driving consumer engagement. Historical views of energy from AMI are just never going to be interesting on an ongoing basis for consumers. The challenge of course is that to get this level of detailed and real-time data, we need to install hardware in consumers home – so has an obvious associated cost. But, as applications like Sense gain consumer adoption (as part of the overall consumer trend towards smarter homes), it is likely that this functionality will over time become a standard feature of a connected home. The core high resolution metering needed by applications like Sense could in the future be built into utility meters or into the core electrical infrastructure of the home.

Ubiquitous real-time device-level energy tracking technology may seem far-fetched. Today, getting a Sense in your home costs hundreds of dollars and most people don't know it exists. Yesterday, the same could have been said for smart thermostats. Today venture capitalists estimate with a straight face that 90% of adults on earth have a phone—and 60% are smart! Even if they are off by a year or two, does it matter? So, in the context of today, it seems plainly clear that the following will be true tomorrow:

- More data about energy will be more available, more meaningful, and more affordable
- More value will be provided, through more use cases, to more entities and people

As a result, cost-effectiveness tests for these technologies will soon be met with only modest growth rates of benefits, and cost reductions through economies of scale and integration.

But, it will take some time—and perhaps a new operating system of energy engagement economics—for energy stakeholders to come to actionable and functional terms with the outcomes of the injection from today's meteoric information technology market ecosystems. So, we can and should do work now to understand the path for using this kind of information and help shape the evolution of this technology to not only match the needs of end consumers, but to allow them to be better consumers of energy through efficiency, load shifting, strategic electrification, and so on.

But it would be short sighted to limit our scope to studying only on how this tech can help homes it is in. It may well be the case that the more "shovel-ready" set of opportunities for this emerging technology could be in the development of robust and representative panels of homes. Constructed so as to be representative of the population at large, or aimed at populations of particular strategic, regulatory, and moral interest (such as low-income, Multi-family, electric vehicles, solar PV, storage, ductless heat pumps etc.) The benefits of developing these "Nielsen" style sample groups could be used to immediately contribute to more detailed usage and savings forecasts for demand resource planning efforts, regulatory scenario modeling, and a host of M&V related activities.

So, is disaggregation perhaps better described not as a single holy grail, but a multitude of mythological objects? With real-time device-level data as the Seven League Boots. They can carry utility programs across big-data chasms and uncanny valleys. Meaningful energy insights could be Excalibur. Once freed from the stone of inscrutability, they cut the Gordian Knot of engagement. Inside is the Philosopher's Stone of authentic customer relationships. Together, the unsustainable lead of today's energy system could be transformed into the eternal life gold of a clean, reliable, and least cost future. Too much? Well maybe it is just the Holy Grail, after all.

# References


Accenture, 2013 *New Energy Consumer Handbook* Accenture Innovation Center for Energy and Utilities www.accenture.com/_acnmedia/Accenture/next-gen/insight-unlocking-value-of-digital-consumer/PDF/Accenture-New-Energy-Consumer-Handbook-2013.pdf

Armel, C. Gupta, Abhay, et. al. 2012 *Is Disaggregation the Holy Grail of Energy Efficiency? The Case of Electricity*, Precourt Energy Efficiency Center, Stanford University, Technical Paper 2012-05-1, 2012

Ashby, K. 2017 *Behavior Insights and Tools: How social Science Has Been—and Could Be—Applied to Connected Programs* Consortium for Energy Efficiency (CEE)

Bertoldi, P., Serrenho, T., and Zangheri, P. 2016 *Consumer Feedback Systems: How Much Energy Saving Will They Deliver and for How Long?* ACEEE Summer Study on Energy Efficiency in Buildings

Delforge, P., Schmidt, L., and Schmidt, S., 2015 *Home Idle Load: Devices Wasting Huge Amounts of Electricity When Not in Active Use* Natural Resources Defense Council Issue Paper.

Dougherty, A., & Provencher, B. (2013, August). Impacts of Feedback Programs: Generating Comparable Impacts across Varying Program Design Models. Presented at the InternationalEnergy Program Evaluation Conference, Chicago, IL.

Fogg, BJ.  *A Behavior Model for Persuasive Design* Proceedings of the 4th International Conference on Persuasive Technology*.* Article No. 40 Stanford University

Kelly, J. and Knottenbelt, W. 2016 *Does disaggregated electricity feedback reduce domestic electricity consumption? A systematic review of the literature* arXiv:1605.00962v2 [cs.CY] 4

Kelly, J. 2017 *Disaggregation of Domestic Smart Meter Energy Data.* PhD Thesis University of London Imperial College of Science, Technology and Medicine Department of Computing

Localytics *Average user Retention by Industry – Three month average* based 12-month analysis of 2016 data. “What’s a good retention rate?” 2017 info.localytics.com/blog/mobile-apps-whats-a-good-retention-rate

M. Zeifman, K. Roth. 2011. *“Nonintrusive Appliance Load Monitoring: Review and Outlook”* IEEE International Conference on Consumer Electronics 2011: 239 - 240