# Differential Privacy for Expanding Access to Building Energy Data

*McGee Young, Recurve*
*Marc-Antoine Paré, Recurve*
*Harry Bergmann, United States Department of Energy*

## Abstract

This paper reviews the current literature on differential privacy, explains its relevance to existing energy data privacy rules, and describes how different permutations of core mathematical principles can be applied to it. We then review a variety of use cases that have been developed as part of a DOE grant that has supported this research and provide insights based on a proof of concept differential privacy library that we have developed. Finally, we discuss technical and policy pathways forward.

## Introduction

Demand side programs play an important role in grid decarbonization scenarios over the next two decades. Given that the carbon-intensity of electricity increasingly depends on when and where it is consumed, planners and program implementers are transitioning from a focus on total savings to one that prioritizes time- and location-sensitive impacts. One key impediment to this transition is the fact that individual building energy consumption is generally considered to be protected under most privacy regulations (Veliz and Grunewald 2018). However, these protections are equivocal and unevenly applied (Ashar et al. 2017). Some rules of thumb have been turned into actual rules, such as aggregation rules that require data releases to roll together multiple buildings in order to conceal the identities of particular consumers. However, in other jurisdictions, individual customer level data is freely available.[1]

The seismic shift in market conditions currently underway raises fresh questions about long-standing firewalls between customers, utilities, and third-parties that prevent broad sharing of customer energy consumption data. While utilities need individual data for billing purposes, significant controversy surrounds data sharing with third parties without express customer content. Proponents of liberalized access argue that as demand side programs consider compensation schemes based on metered savings, outside parties must be able to gain access to

[1] See, e.g., http://gainesville-green.com/

individual building-level data.[2] Similarly, researchers have chafed at the paucity of data available that reflects real-world conditions in buildings.[3]

In this paper, we describe a system for providing privacy-preserving access to otherwise private customer-level energy data. We employ *differential privacy* to provide a mathematically rigorous foundation for privacy protection. We analyze the trade-off between privacy protection and statistical utility for two characteristic use cases - customer targeting and comparison groups. Finally, we describe a system architecture that has been developed by LBNL, NREL, DOE, and Recurve. This system architecture will be launched in a proof of concept demonstration in 2020 with release as a fully functional open-source module planned for 2021.

## Differential Privacy: Background

Differential privacy is a mathematical framework with composable, flexible privacy guarantees that hold regardless of an attacker's existing knowledge about the underlying data. The mechanism underlying this technology was first described in 2006 by Cynthia Dwork and has received extensive academic treatment since then (see e.g., Dwork and Roth 2013). A number of parties have deployed the technology in industry settings, notably a 2018 Uber/Berkeley initiative designed for running arbitrary SQL queries against rider data, in response to GDPR (Johnson et al. 2017).

The basic workings of differential privacy are intuitive to understand and consistent with many existing data aggregation rules. Imagine that you would like to learn what the average annual energy consumption is for a group of buildings, while still preventing identification of any particular building's consumption. Currently in California, for example, if these are residential buildings, at least 100 buildings must be represented in the aggregation with no single building representing more than 15% of the total consumption. (If the buildings belong to the commercial sector, the required number of buildings is relaxed to 15 in the aggregation.) These techniques are designed to add noise to the result in a way that obscures the contribution of any individual building to the final average, even if you knew the consumption of all the other buildings.

Differential privacy adds this noise in a way that meets precise mathematical definitions. Whereas the noise added through conventional aggregation techniques may leave sensitive data exposed in unpredictable ways, the noise added by a differential privacy mechanism provides a different sort of guarantee. Where data has been protected through differential privacy, an attacker with perfect information about a database can only increase their knowledge about whether an individual's data was even included in the dataset up to a threshold $\varepsilon$. That is, the number of queries that can be executed against a database can be limited so that up to a certain

---

[2] One example organization pushing for liberalizing data access is Mission Data. http://www.missiondata.io

[3] A recent report documenting efforts to use data to improve energy efficiency can be found at https://www.imt.org/resources/putting-data-to-work-how-cities-are-using-building-energy-data-to-drive-eff/

limit, it will be impossible to know whether or not any particular building's energy use is included within the dataset in the first place. Once this threshold is crossed, the attacker would know whether or not an individual building's data was included, but the data would still be aggregated and subject to the same protections as existing aggregation rules, meaning that even if it were known whether or not an individual's data was represented, the individual's contribution to the aggregation would still not be known. In short, differential privacy adds a second layer of protection to ensure anonymity of individual user data within an existing set of dataset aggregation protections.

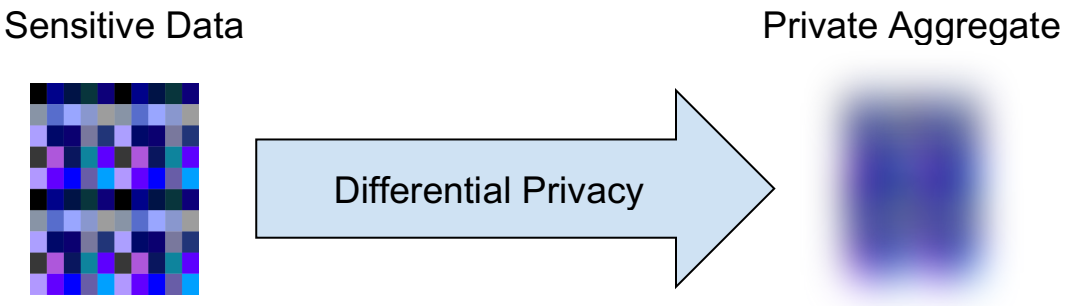


Figure 1. Differential Privacy aggregates data and adds noise to make it difficult to determine the contribution of any particular element in the dataset.

Mathematically, $\varepsilon$ is defined as the ratio of probabilities between two databases that differ by a single element. This ratio must hold for all possible outputs and all possible pairs of neighboring databases. If an attacker were to try to infiltrate a dataset by utilizing other available data (such as public records), the $\varepsilon$ protection would still hold for the differentially private dataset. As Figure 1 shows, the differentiation between query outputs is accomplished by varying the probability of the validity of any given output. One way to think about this is to imagine that you would ask someone a yes/no question and there would be a varying probability that the answer they provided was true or not. Only by repeatedly asking this question could you start to guess whether the answers were truthful. The $\varepsilon$ value prevents you from asking this question enough times to develop any certainty about the responses.

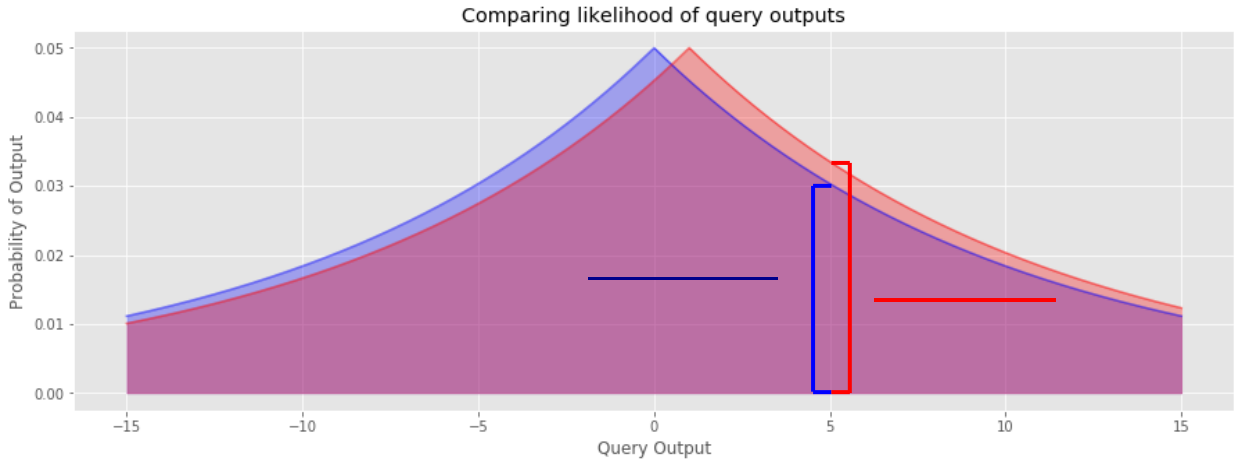


Figure 2. The ratio of the likelihood of any output from a differentially private mechanism with and without any particular element should be bounded by $e^{\varepsilon}$.

A fundamental concept in Differential Privacy is the "privacy budget". Every anonymization technique, from k-anonymity, to the 15/15 rule, to Differential Privacy, reveals some amount of information about the individuals in the dataset. For example, imagine if one were able to find out the consumption information for 14 of the 15 individuals in a dataset anonymized by the 15/15 rule; it would be possible to then deduce that final participant's energy usage. This sort of attack is not just theoretical; Netflix's de-identified movie review database was re-identified using publicly available IMDB data. Even more shocking was researchers' ability to uncover the health information of Massachusetts Governor William Weld. The information had been de-identified in accordance with HIPAA standards, but was re-identified using cross-referenced public voting records.

In contrast to other privacy techniques, Differential Privacy quantifies the risk to each individual contained in a dataset *no matter how much additional data is released about them*. This risk is quantified as $\varepsilon$, the privacy parameter.

Given the vulnerabilities of existing privacy practices, the guarantees offered by differential privacy mechanisms have a number of highly desirable properties.

- They are composable: it is possible to compute an exact bound on privacy loss from multiple statistical releases. Multiple queries against a database gracefully degrade privacy guarantees instead of catastrophically failing (non-binary outcomes are possible).
- The privacy protections are not dependent on an attacker's existing level of knowledge.
- The outputs of a differentially private mechanism are robust to post-processing, not compromising privacy no matter how much additional computation is performed on them.

## Architecture for a Differentially Private Query System

A differential privacy software architecture could take many forms. For energy data, the most common existing use case is an interactive query architecture (see, e.g., https://openei.org/).

We use this typical scenario to describe a simplified data governance structure, shown in Figure 2, whereby a sensitive dataset is managed by a data administrator and users issue queries via a secure interface. This framework provides visibility into how a privacy budget is allocated across an entire dataset that is shared by all users of the system. As queries are issued, the privacy budget is consumed, and, once exceeded, all access to the system is revoked.

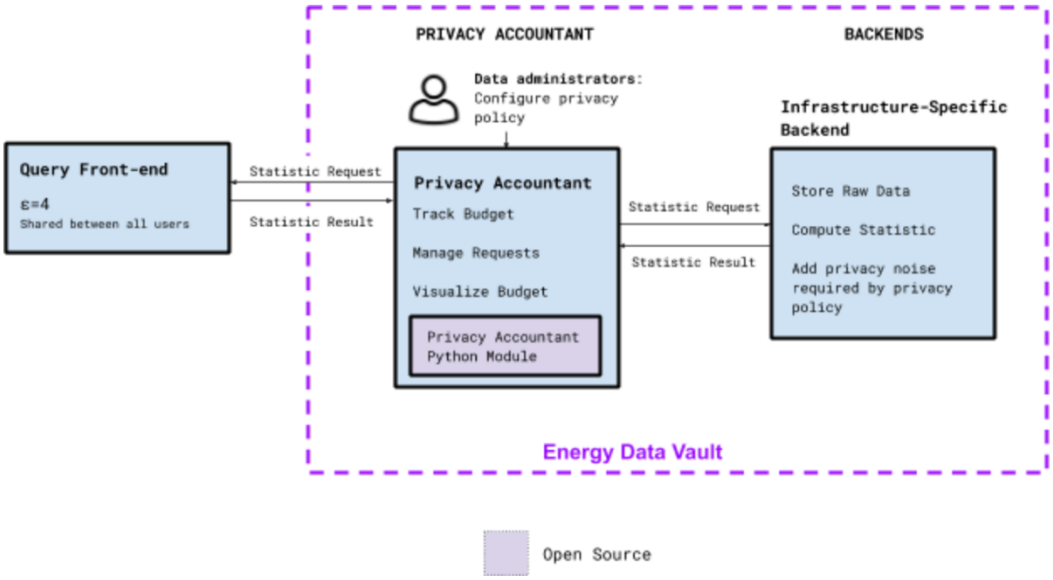


Figure 3. Differential privacy data governance structure.

## Differential Privacy in Action: An Example for Building Energy Consumption

Following the architecture described above, we provide a simple use case that would benefit from a layer of differential privacy protection. Imagine that an agency would like to release the average annual energy consumption for a group of buildings. These consumption values would range from zero to some maximum, $C_{max}$.

Our differential privacy algorithm must compute the amount of noise that will obscure an individual building's contribution to the final average. The basis of this computation is the condition that the maximum amount that a single building can contribute to the average of the group, $\Delta A_{max}$, is

$$\Delta A_{max} = C_{max} / \mathrm{N}$$

where $N$ is the number of buildings in the query.

One common technique for generating noise is the Laplace mechanism, shown in Figure 3, which adds noise derived from a Laplace distribution. A Laplace distribution uses $\varepsilon$, which is the "privacy parameter" to determine the amount of noise that gets added to an aggregation result. For any given desired $\varepsilon$, the Laplace distribution with mean zero and scale $b = \Delta A_{max} / \varepsilon$

quantifies the strength of the privacy guarantee. Counterintuitively, the higher the $\varepsilon$, the less the data is obscured. So a dataset with $\varepsilon$ of 1.0 will be obscured less than a dataset with $\varepsilon$ of 0.1.

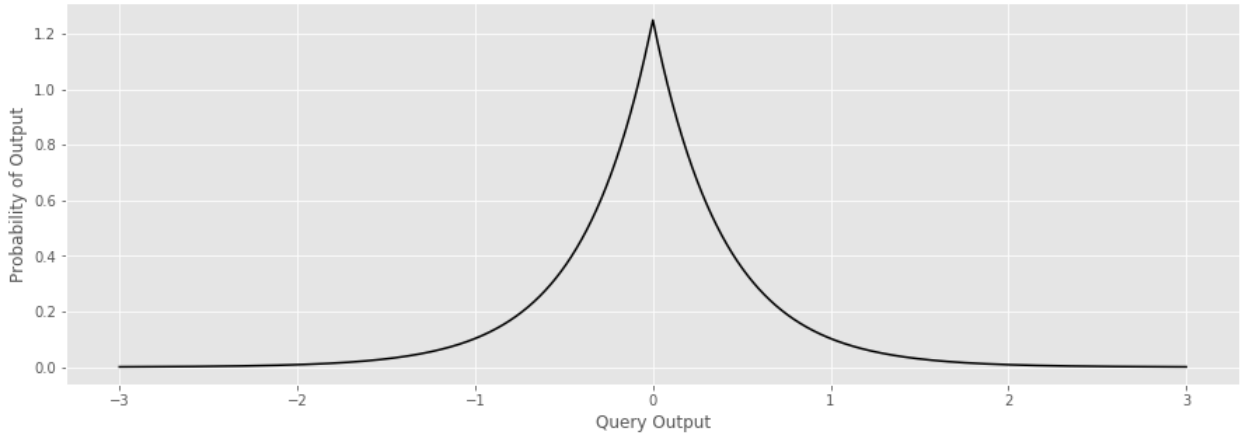


Figure 4. Example Laplace distribution.

Using this probability distribution function, many of the noise values generated are around zero, but there is a reasonable chance that values further away are chosen. In our example, if we had a portfolio of 5,000 homes ($N$=5,000), whose annual consumption ranged from 0 to 30,000 kWh, and a privacy parameter of $\varepsilon$ = 1.0, the final output would be a value sampled around the true mean at ±36 kWh at 95% confidence, as shown in Figure 4.

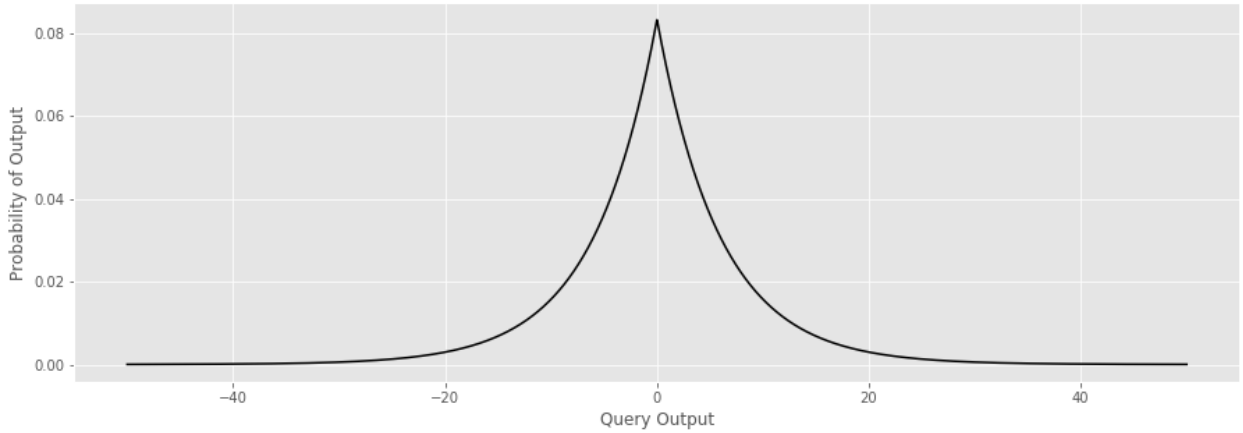


Figure 5. Distribution of differentially private outputs from example query.

## Choosing the Privacy Parameter $\varepsilon$

Core to the differential privacy formulation is the privacy parameter $\varepsilon$ that dictates how much noise is added to output statistics. Unfortunately, the inventors of differential privacy point out that the choice of this parameter is "essentially a social question" left to implementers (Dwork 2008). Existing work provides some guidance as to the range of choices of $\varepsilon$ that should be expected as policy evolves. The Census Bureau arrived at an $\varepsilon$ of 8.9 for the OnTheMap

application (Andersson et al. 2009). Hsu et al. (2014) summarize a number of sources, finding most values to range between 0 and 1, with some as high as 10.

To provide guidance on the necessary privacy parameter for useful outputs, the rest of this paper examines a range of ε choices from 0 to 9 for different use cases. In future work, we hope to provide guidance for data owners to choose privacy parameters based on their legal restrictions, dataset characteristics, and threat models.

## Exploration of Key Use Cases

An initial technical advisory group was created to support the development of use cases for this project. As we describe below, these use cases flowed from existing challenges facing researchers and program administrators as a result of data privacy restrictions. We examine how program administrators might use differential privacy for energy savings calculations from non-participants and how differential privacy could be used by utilities to provide more targeted incentives for efficiency initiatives.

### Energy Savings Calculations in a Comparison Group

Open-source methods such as CalTRACK allow energy efficiency savings to be valued using observed changes in energy consumption. In estimating the impacts of programs, it is often necessary to account for factors that may cause energy use to change, but are exogenous to the program. A very timely example is the impact of Shelter in Place orders related to COVID-19. Many commercial facilities have closed or significantly reduced their consumption, while residential buildings are increasing their energy consumption. To accurately estimate the impact of an energy efficiency program, it would be necessary to filter out the changes due to this large scale exogenous event.

Retrospective program analyses have used the consumption data of similar non-participants to estimate the impact of exogenous factors. But for the purposes of retrospective program analysis, this non-participant data is tightly controlled and available only to small teams of hand selected analysts. For today's market participants that are being held accountable for meter-based savings, real-time access to non-participant comparison group data is essential, especially in the face of unprecedented uncertainty associated with COVID-19. Here is an example of how a differential privacy approach can meet participants' need for comparison group data while retaining customer privacy:

A query may look like the following: compute the average percent NMEC (Normalized Metered Energy Consumption) savings for a portfolio of 2,000 buildings that underwent efficiency projects and a similar portfolio of buildings that did not participate in a program.

NMEC savings are a percentage value, clipped to range from -20% to 20% savings (larger values are generally dropped from the analysis as outliers). The final average of a portfolio tends to be between -4% and 4%, with model error ±1% .

The plot below in Figure 5 shows the error at 95% confidence for this example using the Laplace Mechanism for $\varepsilon$ values ranging from 0.0 to 4.0. The error values here are for percentage points. For example, for an output mean of 3% and $\varepsilon$=0.25, the final output would be 3% ±0.5%.

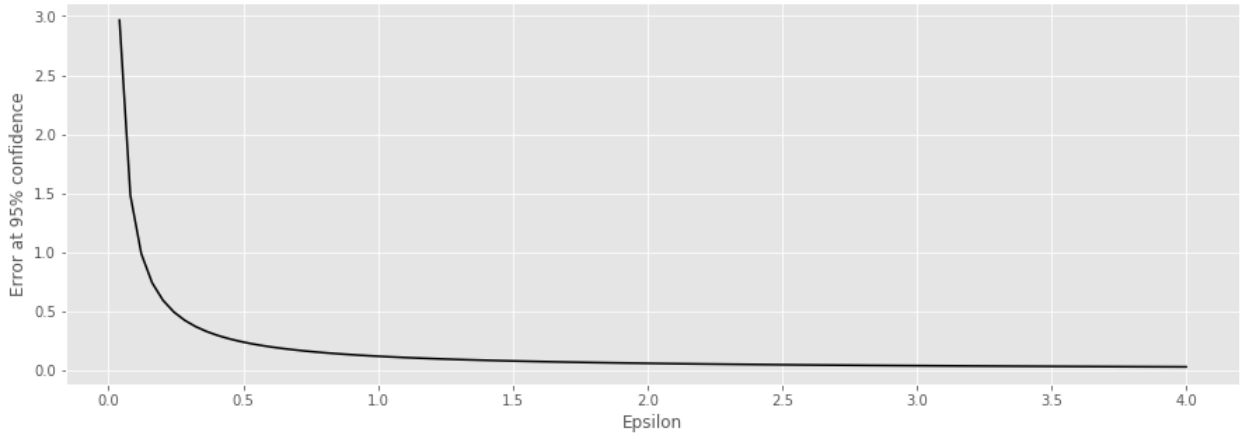


Figure 6. Range of error at multiple $\varepsilon$ values.

Queries with $\varepsilon$ values 0.2 and greater have equal to or less error than the underlying model error. Depending on the use case, these errors are acceptable, though, in general, as little error as possible is desirable.

The energy savings calculation use case also demonstrates the sensitivity of accuracy from query characteristics. A query over a larger portfolio will be more accurate because the scale parameter is reduced proportionally with an increase in $N$. Similarly, if the sensitivity ($\Delta$) of the dataset is decreased by clipping more aggressively, from a range of [-20%, 20%] instead of [-40%, 40%], the statistics output will have less noise introduced by the Laplace Mechanism. The influence of some of these parameter choices is plotted in Figure 6 below:

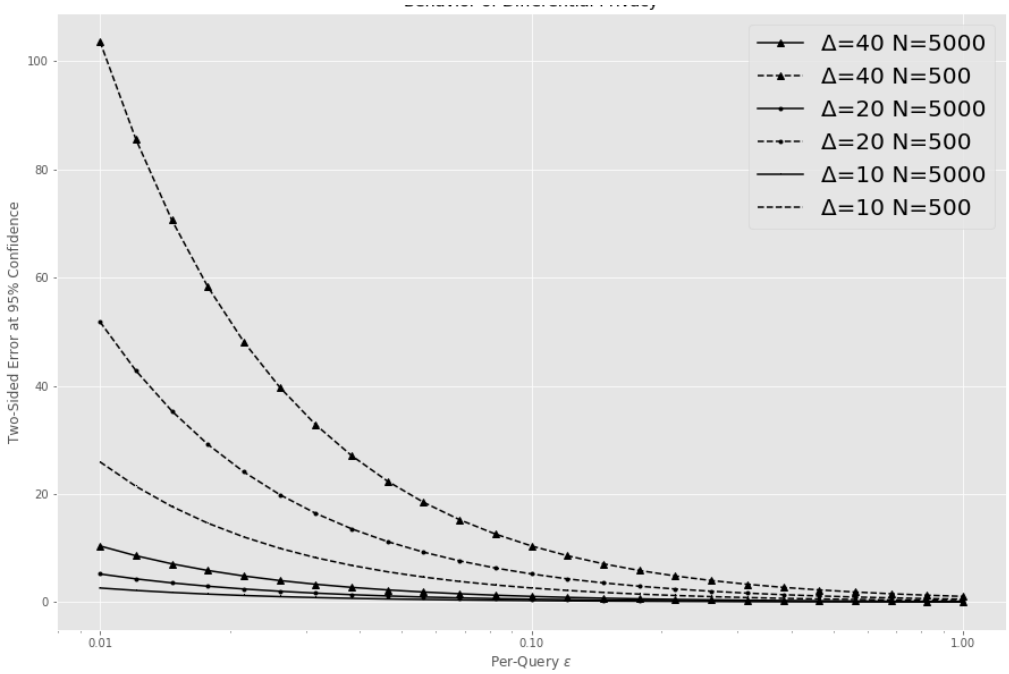


Figure 7. Influence of sensitivity and portfolio size on error.

*Limitations*

It is clear from the plots of output error bounds that differential privacy is not suitable for queries for portfolios on the order of 10s to 100s of sites because of large amounts of output error. In some cases, it might be possible to issue a single query at high $\varepsilon$ to get a reasonable answer for smaller portfolios, but this bumps into a further limitation of the behavior of differential privacy under multiple queries.

The major limitation facing differential privacy for the energy savings calculation use case is managing privacy budgets for multiple queries. The examples so far have considered issuing a single query. In practice, an interactive query system would need to support a large number of queries. With the query model described so far, the privacy impact of multiple queries is computed via composition. There are a number of composition approaches, the simplest formulation being "basic composition" in which the $\varepsilon$ values add up from each of the queries issued to a total $\varepsilon$ value for the entire sequence. In other words, five queries of $\varepsilon$=0.1 have a total $\varepsilon$=0.5. This privacy impact adds up quickly, exhausting even the most generous overall privacy budgets after tens to hundreds of queries.

While differential privacy for energy savings calculations faces some limitations, there are ways forward. The results so far show that, for high value queries, differential privacy offers strong privacy protection and reasonable accuracy. To expand the number of queries supported by a differential privacy system, there are alternative composition formulations like Advanced Composition and Zero-Concentrated Differential Privacy that allow for more queries at the expense of slightly more complicated privacy models. The accuracy and composition problems can also be addressed by incorporating a contractual layer to the system that limits sharing of

outputs to specified parties, which would allow more generous privacy budgets since the budget did not have to be shared globally.

**Targeting**

The second use case that we profile shifts the privacy logic in a different direction. As many program evaluations have shown, actual savings from energy efficiency measures vary quite significantly from building to building. Recurve's Resource Planning tools have shown that existing patterns of energy use within a building can be highly predictive of the potential for future energy savings in a building (see also, Borgeson, 2018). Being able to more accurately target buildings with a high potential for valuable demand flexibility gives third parties an enormous advantage, especially for utility procurements in which incentives are tied to the value of the savings at particular times of day. Utilities would like to be able to provide third parties with portfolios of candidate sites for retrofits that greatly increase the likelihood of success without revealing any individual site's energy consumption patterns.

For this use case, a differentially private targeting query can be constructed using a randomized response method. This method, developed in the 1960s for psychological research (Fox 1986), is at the core of Google's differentially privacy system, RAPPOR (Erlingsson et. al 2014).

To understand how a randomized response privacy protection works, consider how you might obscure answers to a yes/no question. In this case, participants are asked a yes-or-no question and flip two coins secretly. If face-up, the participants respond truthfully to the query. If face-down, the participant responds with the value shown by the second coin, i.e. randomly. With a large enough pool of participants, the effect of random response can be averaged out of the final out of "yes" and "no" responses.

Randomized response provides differential privacy. For a fair coin ($p_{\text{face-up}} = 0.5$), the equivalent $\varepsilon$ is equal to ln(3). By changing $p_{\text{face-up}}$, the privacy guarantee and accuracy can be adjusted.

This procedure can be adapted to the targeting use case. As a simplified example, consider a program that would like to target the top 50% of energy consumers in a population. A database is prepared with an entry for each building: one if it is in the top 50% of energy consumers and zero otherwise. The randomized response procedure is performed for each building and the set of all buildings that responded positively is returned.

Some fraction of these positive responses will not actually fall in the top 50% of energy consumers because of the randomized response procedure. Also, some fraction of the participants that were in the top 50% will not be included in this output set. The following plots in Figures 7 and 8 show the effect of various choices of the privacy parameter on the fraction of sites correctly targeted:

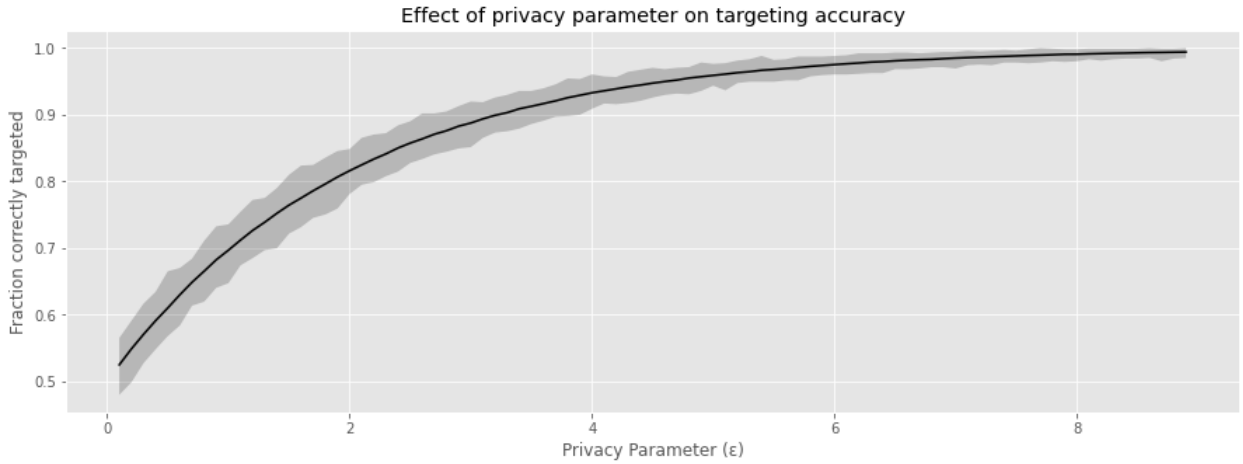


Figure 8. Effect of privacy parameter on targeting accuracy. 500 trials run at each ε.

For example, at an $\varepsilon$=2, the fraction of correctly targeted sites is about 0.82. A portfolio with 100 sites sampled from the targeted group will, on average, have 82 sites correctly targeted and 18 incorrectly targeted.

Interestingly, this method works equally well irrespective of targeting rigor. That is, changing the fraction of the population that meets targeting criteria does not have an effect on the accuracy of targeted portfolio:

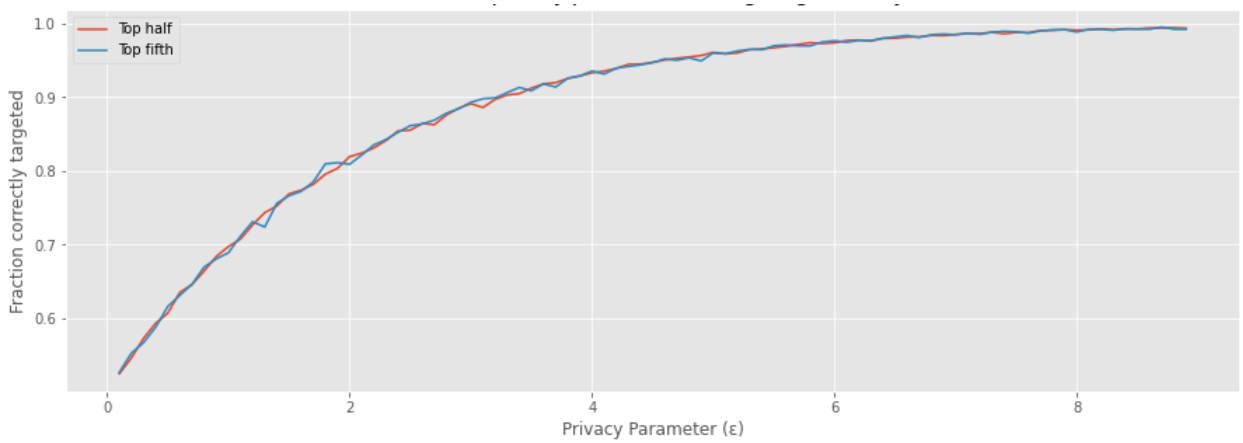


Figure 9. Comparison of impact of targeting rigor on accuracy.

*Limitations*

Like the comparison group use case, targeting requires relatively high values of $\varepsilon$ for usable results. The privacy parameter values should be greater than 1.0 to correctly target at least 70% of sites. For the targeting use case, this may not be problematic either, as these queries only have to be run once against sensitive data to inform a multi-year program.

Another limitation of targeting via randomized response is that a target portfolio will always have false positives. It is possible that this will cause some confusion among stakeholders that the target portfolio is only mostly full of high value targets.

## Conclusion

We have shown how differential privacy can be applied to valuable use cases in the energy efficiency sector. This work is currently undergoing pilot testing and will continue to evolve. In many ways, the principle of differential privacy parallels the emergence of pay-for-performance in the energy efficiency sector. Privacy regulation in this sector, namely the 15/15 standard, has relied on ad-hoc definitions without quantitative rationale for decision-making. It is exceedingly difficult to apply this standard to new analytics needs or understand the impact of data releases over time. Managing energy data privacy within the status quo of policy and technology is like managing the effectiveness of large-scale energy efficiency interventions without measuring performance at the meter; it works when everyone agrees based on common sense, but breaks down under numerical scrutiny, leaving potential benefits unrealized. In this sense, the benefit of differential privacy isn't as a magic black box that grants privacy for free, but as a principled approach that allows decision-makers to quantitatively understand their privacy decisions. This firm foundation can potentially unlock a great variety of use cases while simultaneously providing stronger privacy guarantees.

## References


Andersson, F., Abowd, J., Graham M., Wu J., and Vilhuber, L., 2009. Formal Privacy Guarantees and Analytical Validity of OnTheMap Public-use Data. In Joint NSF-Census-IRS Workshop on Synthetic Data and Confidentiality Protection. Cornell University, Suitland, MD. https://ecommons.cornell.edu/handle/1813/47672

Asghar, M., G. Dán, D. Miorandi and I. Chlamtac, "Smart Meter Data Privacy: A Survey," in IEEE Communications Surveys & Tutorials, vol. 19, no. 4, pp. 2820-2835, Fourthquarter 2017, doi: 10.1109/COMST.2017.2720195.

Borgeson, Sam, et. al., Energy Efficiency Program Targeting: Using AMI data analysis to improve at-the-meter savings for small and medium businesses." CALMAC Study ID PGE0421.01.

Dwork, C., 2008. Differential Privacy: A Survey of Results, in: Agrawal, M., Du, D., Duan, Z., Li, A. (Eds.), Theory and Applications of Models of Computation. Springer Berlin Heidelberg, Berlin, Heidelberg, pp. 1–19. https://doi.org/10.1007/978-3-540-79228-4_1

Dwork, C., Roth, A., 2013. The Algorithmic Foundations of Differential Privacy. FNT in Theoretical Computer Science 9, 211–407. https://doi.org/10.1561/0400000042

Erlingsson, Ú., Pihur, V., Korolova, A., 2014. RAPPOR: Randomized Aggregatable Privacy-Preserving Ordinal Response. Proceedings of the 2014 ACM SIGSAC Conference on Computer and Communications Security - CCS '14 1054–1067. https://doi.org/10.1145/2660267.2660348

Fox, J.A, and Tracy, P.E 1986, *Randomized response*, Quantitative applications in the social sciences, SAGE Publications, Inc., Newbury Park, California, [Accessed 13 March 2020], doi: 10.4135/9781412985581.

Hsu, J., Gaboardi, M., Haeberlen, A., Khanna, S., Narayan, A., Pierce, B.C., Roth, A., 2014. Differential Privacy: An Economic Method for Choosing Epsilon, in: 2014 IEEE 27th Computer Security Foundations Symposium. Presented at the 2014 IEEE 27th Computer Security Foundations Symposium (CSF), IEEE, Vienna, pp. 398–410. https://doi.org/10.1109/CSF.2014.35

Johnson, N., Near, J.P., Song, D., 2017. Towards Practical Differential Privacy for SQL Queries. arXiv:1706.09479 [cs]. https://doi.org/10.1145/3177732.317773

Véliz, C., Grunewald, P. Protecting data privacy is key to a smart energy future. *Nat Energy* 3, 702–704 (2018). https://doi.org/10.1038/s41560-018-0203-3