Getting Real with Energy Data: Using the Buildings Performance Database to Support Data-Driven Analyses and Decision-Making

Richard E. Brown, Travis Walter, Laurel N. Dunn, Claudine Y. Custodio, and Paul A. Mathew, Lawrence Berkeley National Laboratory
D. Magnus Cheifetz, Building Energy Inc.
Elena Alschuler, U.S. Department of Energy
Jessica Knapstein, Energetics Inc.

ABSTRACT

The lack of empirical data on the energy performance of buildings is a key barrier to accelerating the energy efficiency retrofit market. The DOE's Buildings Performance Database (BPD) helps address this gap by allowing users to perform exploratory analyses on an anonymous dataset of hundreds of thousands of commercial and residential buildings. These analyses enable market actors to assess energy efficiency opportunities, forecast project performance, and quantify performance risk using empirical building data. In this paper, we describe the process of collecting and preparing data for the database, and present a peer-group analysis tool that allows users to analyze building performance for narrowly defined subsets of the database, or peer groups. We use this tool to explore a case study of a multifamily portfolio owner comparing his buildings' performance to the peer group of multifamily buildings in the local metro area. We also present a performance comparison tool that uses statistical methods to estimate the expected change in energy performance due to changes in building-component technologies. We demonstrate a low-effort retrofit analysis, providing a probabilistic estimate of energy savings for a sample building retrofit. The key advantages of this approach compared to conventional engineering models are that it provides probabilistic risk analysis based on actual measured data and can significantly reduce transaction costs for predicting savings across a portfolio.

Introduction and Overview

The DOE's Buildings Performance Database (BPD) is a decision-support platform, comprised of a database and data analysis tools, that enables statistical analysis of building energy performance, operational, and physical characteristic data. The database contains measured data for over 750,000 existing commercial and residential buildings, allows analysis without revealing sensitive information about individual buildings, and is put through a stringent cleansing process to enable confidence in analysis. The feature that distinguishes the BPD from other similar databases is the ability to analyze narrowly defined peer groups—such as market regions, and/or specific building or equipment types—at a localized or factor-specific level appropriate for decision-makers such as building owners and policymakers. The goal of the BPD is to enable market actors to assess energy efficiency opportunities, forecast project performance, and quantify performance risk using empirical building data.

There is a common saying that "All real estate is local." Therefore, the goal of the BPD is to leverage the law of large numbers to achieve validity at the local level, not to achieve a representative national sample. Thus, the BPD is a complementary effort to the Residential and Commercial Building Energy Consumption Surveys (RECS and CBECS). CBECS and RECS are statistically valid samples of all commercial and residential buildings in the country, and are in fact two of the data sources for the BPD. However, CBECS and RECS are not valid samples for narrowly defined peer groups. The BPD might contain enough data to approximate the underlying distribution for specific local markets and building types. For example, the BPD contains data collected under energy use disclosure ordinances in San Francisco, Seattle, Washington, D.C. and New York City. Since compliance with the ordinances in these cities is relatively high, the BPD likely contains a nearly complete sample of the population of buildings to which the ordinance applies.

The two analysis tools currently available in the BPD, the Peer Group Tool and Performance Comparison Tool, allow users to compare performance trends among similar buildings, identify and prioritize cost-saving energy efficiency improvements, and assess the range of likely savings from these improvements. While the BPD platform offers various analysis tools through a public web site, third parties can also access the database through an Application Programming Interface (API). API users can query the database to conduct their own analysis without compromising the security or anonymity of the database. The API can be used for app development, visualization or research purposes, or to integrate the BPD with privately held data.

The BPD provides a secure and anonymous means of publishing building data originally collected for other purposes, including energy efficiency incentive programs, academic research, efficiency projects completed by owners or contractors, and energy disclosure laws. Once public, these data may influence local real estate markets by revealing trends and opportunities. For example, determining the relative energy performance of a building may influence what prospective buyers or tenants are willing to pay for those buildings. The BPD also provides a simple way for building owners and managers to benchmark buildings against a user defined peer group enabling them to recognize high or low performing buildings within their portfolios, prioritize buildings for investment, and identify regionally specific improvements with a demonstrated savings impact.

The BPD can also inform policymakers, utilities and energy efficiency program administrators by providing information about the local building stock. The Performance Comparison Tool can help these stakeholders to identify building types and energy efficiency measures that have the greatest energy savings potential, or measures with low uncertainty in savings estimates. By utilizing local building stock and energy conservation measure performance data, program administrators can also improve policy and program targeting and design to maximize savings potential. For example, policymakers can analyze subsets of the local building stock to determine whether specific equipment types are correlated with high energy consumption in their region, and apply these results to tailor energy efficiency incentive programs.

Finally, the measured data for real buildings stored in the BPD can satisfy a deficiency in the energy efficiency investment community by providing more detailed information on local real estate markets and enabling quantitative risk analysis. Currently, the market conventionally assesses energy efficiency project performance through return, simple payback, or net present value alone with no investment risk analysis. Financiers require an understanding of the range of savings likely to be achieved by performing the same measures over a large group of similar buildings in order to recognize energy efficiency lending as a unique asset class and provide capital at scale.

About the Data

The BPD contains data from over 25 source data sets, totaling more than 750,000 buildings (as of May 2014), submitted voluntarily by both public and private data contributors. There is no upper bound on the amount of data the database can hold. While the dataset contains primarily residential records—about 90% of the buildings in the database—there are many more residential than commercial buildings in the U.S.. Comparing the number of residential and commercial building records in the BPD, we find that the database contains 0.9% of the U.S. commercial building stock (EIA 2003), and 0.6% of the U.S. residential building stock (EIA 2009). Primary data contributors include federal government agencies such as the Environmental Protection Agency (EPA) and the Department of Housing and Urban Development (HUD), state and local governments, large portfolio owners, energy efficiency programs and private companies (DOE 2014a).

All building records contain basic building characteristics such as location, building type, and energy use; these fields constitute the BPD's minimum data requirements. These data are useful for analyses such as: comparing a building or portfolio performance to similar buildings in the BPD, or identifying high or low performing buildings in a peer group. The BPD currently analyzes building records as "snapshots" at a point in time. Changes in building characteristics, for example due to a retrofit, are captured by entering the building as a new record with modified characteristics. As the dataset grows, the BPD will allow users to analyze buildings longitudinally, capturing pre- and post-retrofit asset information and energy use data. Although the BPD has few mandatory fields, the database accepts records that include upwards of 350 data fields, as specified in the Building Energy Data Exchange Specification (BEDES) (DOE 2014b). Additional fields include detailed characteristics about building systems, operations, and characteristics. The BPD schema can also hold high-resolution energy use data collected at daily, hourly, or 15 minute intervals.

Interval meter data and equipment or asset level data are anticipated to become increasingly prevalent as data-driven products and services gain traction in the market. As stakeholders begin to collect and contribute asset rich data, the analysis capabilities of the BPD will continually increase; the breadth and depth of analysis is only bounded by the amount of data in the database. Existing tools that collect and analyze building data are the EPA's Portfolio Manager, and the DOE's Home Energy Score and Commercial Building Energy Asset Score (Asset Score). Portfolio Manager tracks building energy and water consumption, and computes weather normalized energy use intensity (EUI) and a whole-building score based on actual energy performance. The Home Energy and Commercial Energy Asset Scores collect information about building materials, equipment, and other physical characteristics, to provide a whole-building score and identify inefficient systems and potential capital upgrades based on as-built physical characteristics (independent of operations). Together, these tools provide a complete picture of a building's physical and operational characteristics, which can then be contributed to the BPD.

Data contributed to the BPD comes from many sources, in many different formats and with variable data quality. Before being imported to the database, data undergoes semi-automated data preparation process, illustrated schematically in Figure **1**. The purpose of the data preparation process is to maintain data quality within the BPD and to ensure that all data enables meaningful and confident analysis; no data is entered directly into the BPD without first going through the data preparation process.

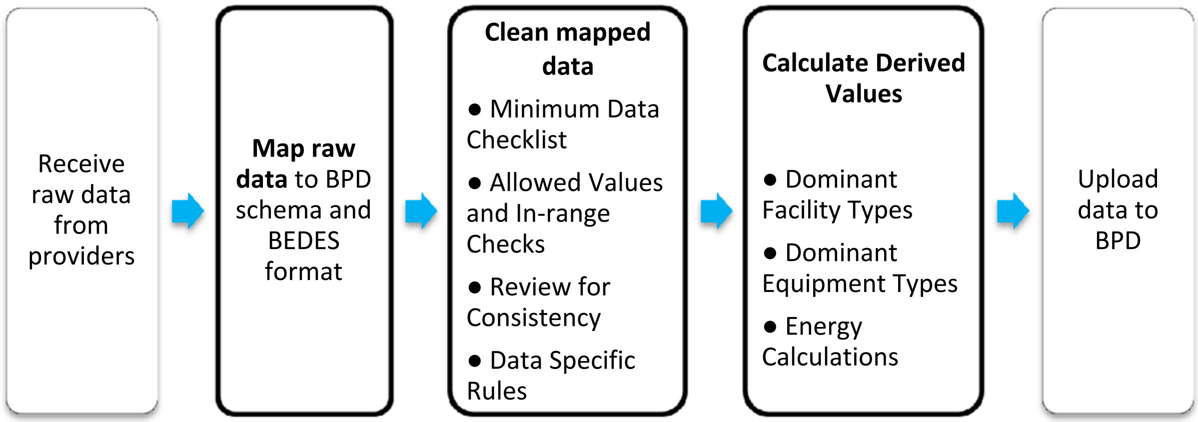

Figure 1: Schema detailing the BPD data preparation process including: mapping data fields, cleansing data entries, verifying consistency within data entries, and deriving values (e.g. energy totals and dominant asset types).

As most datasets adhere to their own schema or data format, raw data from data providers is first mapped. This process involves translating field names and terminology used by each data contributor to equivalent terminology specified in BEDES and encoding it to the BPD data schema. Mapping decisions are made by LBNL, with data dictionaries (if available) or correspondence with data contributors providing clarification as needed. Translating data into a consistent format and language facilitates computational analysis and aggregation into the database.

Data cleansing involves an automated process of confirming 1) that buildings meet minimum data requirements, 2) that data entries make sense, and 3) that building records are internally consistent. A document detailing these processes is in review, and will be made available on the DOE's BPD website (DOE 2014a). In the first stage of cleansing, buildings are removed from the dataset if they contain missing, erroneous, or otherwise suspect data in required data fields. During the second stage, ranges of data values are checked using rules informed by building science, and realistic values for each field. Finally, each building record is reviewed for internal consistency and inconsistencies addressed through deletion of the building or of a particular field. The cleansed dataset is then spot-checked for any other abnormalities by the data processer. Data cleansing is an iterative process; erroneous data are removed and the cleansed data is put through the cleansing process again.

To confirm consistency within a building record, value thresholds are set using other values reported for the same building. For example, if the heated floor area of a building exceeds its gross floor area, the entry for heated floor area is deleted. Any indication of "bad data" in required data fields, such as gross floor area, result in deletion of the entire building record. In cases where inconsistencies are identified between a required and an optional field, as in the example above, the required field is assumed to be accurate. While this protocol may seem somewhat arbitrary, the BPD required fields are the more commonly reported types of information about buildings, thus more likely to be properly collected and quality controlled. Once data is cleansed, fields displayed in the online application are calculated using information provided in raw data fields. Calculated fields include [1] annual energy use (including site, source, electricity and fuel) for the most recent year of continuous data, [2] the measurement year corresponding to the energy use interval, [3] dominant building type, and [4] dominant building assets.

As the database increases in size, it becomes more likely that buildings will be duplicated in the database. To avoid duplication, the cleansing methods are being updated to identify buildings in a new dataset that are already in the database. As an example, any ENERGY STAR Labeled building located in San Francisco is likely to be included in datasets contributed by both ENERGY STAR and the City of San Francisco. When new datasets are added, the database will be scanned for duplicates to avoid treating one building reported twice as two distinct buildings. The algorithm for identifying duplicates will flag buildings that share the following properties within some tolerance: [1] postal code, [2] floor area, [3] whole building energy use, and [4] energy measurement year. Buildings below a certain floor area will be excluded from screening because homes or apartments that are similar in construction may use roughly the same amount of energy.

Peer Group Analysis Tool

One application of the BPD is to analyze the performance of narrowly defined subsets of the database, or building peer groups. Peer groups are selected by filtering the database based on building characteristics, such as building type, location, size, vintage, or heating and cooling equipment. Filtering effectively allows users to choose the amount of diversity among buildings within the peer group by either narrowing or broadening the filters that define the peer group. While narrowing filters reduces the size of a peer group, the number of buildings in data-rich subsets of the BPD affords users the option to select very specific database filters and still yield sizeable peer groups.

Peer group analysis can be used to compare performance of a building (or portfolio of buildings) to other similar buildings in the building stock. The results of these analyses can also help users to identify differences in energy consumption or characteristics between different subgroups within a peer group. Summary statistics of a peer group's characteristics, such as energy use intensity or floor area, can also provide actionable metrics for evaluating building performance. To illustrate the analysis capabilities of peer group analysis, we consider three different applications of the tool. These include identifying 1) low-performing buildings within a peer group and 2) characteristics shared among low-performing buildings.

In the first example, a user can filter the database to isolate a peer group in the BPD that is similar to a building or buildings of interest to the user. If we imagine a real estate company that owns 10 multifamily residential buildings in New York City, the portfolio's peer group in BPD contains more than 5,800 buildings. Energy use for the peer group can be visualized using a scatterplot, as in Figure **2**, where the median energy use intensity for the peer group is denoted by a blue circle—in this case 78 kBtu/ft^2-year—and quartiles are denoted by the background shading. The summary statistics of a peer group can be used to develop metrics for identifying high- or low-performing buildings. For example, the portfolio owner can plot one of its own buildings on the graph (denoted by the orange circle) and see how its performance compares to that of its peers. In this example, the building performs in the bottom quartile (denoted by the red shaded region at the top of Figure **2**) of its peer group. These analysis results can be used to identify buildings that are optimal candidates for energy audits or efficiency upgrades.

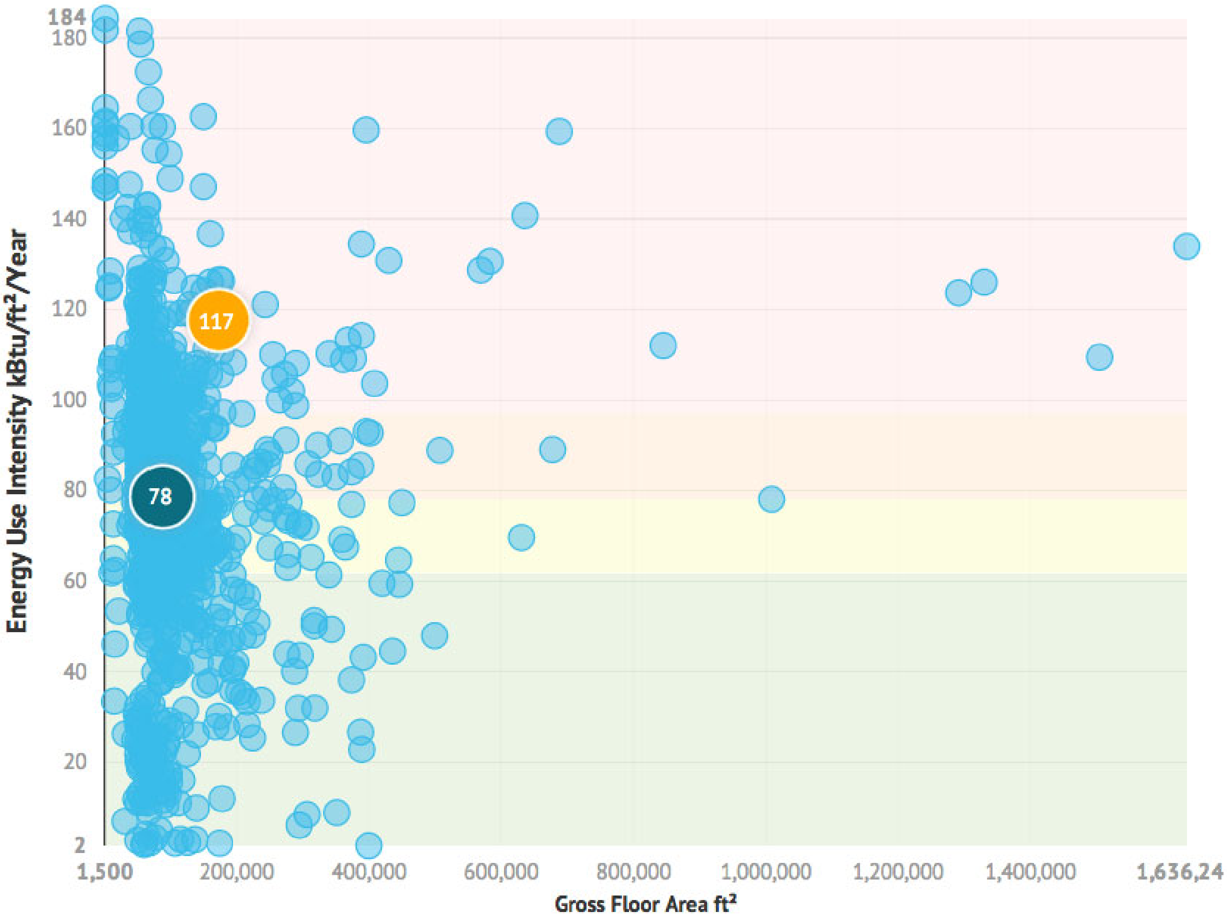

Figure 2: Scatterplot of energy use intensity versus floor area for a random sample of 5,802 multifamily residential buildings in New York City in the BPD with median energy use intensity equal to 78 kBtu/ft^2-year (denoted by the blue circle), and an example user building (denoted by the orange circle). (Source: Building Energy, Inc.)

A second application of peer group analysis is to identify the subsets within a peer group with the lowest performance. Say, for example, we were to filter the peer group above by postal code, comparing the distribution of energy use for buildings located in different regions within the city. This analysis could show that buildings in a certain neighborhood tend to use more energy than buildings elsewhere in the city. The results of this analysis could be used to target energy efficiency programs towards a specific subset of the building stock. This relatively low-effort peer group analysis can inform policy design to maximize impact on energy use.

In order for peer group analyses to yield actionable results, buildings in a BPD peer group, or sample, must be representative of the underlying building stock, or population. In this context, the *population refers not to a national or local building stock, but to all buildings in the stock that meet the filter criteria.* If the BPD contains a representative sample of the underlying population, then summary statistics of the sample—such as the mean, median, and quartiles—will be roughly equivalent to summary statistics of the population. Means for testing representativeness are under development, and discussed further in section 6.

Performance Comparison Tool

Another application of the BPD is to apply data-driven algorithms to understand the physical and operational drivers of differences in building performance. In this section, we describe an analysis driven by BPD data that uses statistical methods to estimate the expected change in energy performance due to changes in building-component technologies. The methods described here have been developed as the analysis approach for the BPD's Performance Comparison Tool.

We construct a multivariate linear regression model that predicts EUI using building characteristics. The model includes several predictors, some numeric (e.g., operating hours, number of occupants, wall R-value) and some categorical (e.g., facility type, climate zone, lighting type). In some cases, the model uses transformations of the predictors (e.g., log(floor area) rather than simply floor area). In addition, the model uses cross terms that capture interactive effects between two different predictors (e.g., heating type and climate zone, cooling type and lighting type).

While the BPD contains over 750,000 buildings, most of these buildings do not report asset data. Depending on the amount and type of asset data reported by the building in the peer group, we decide algorithmically which predictors to include in the model so as to maximize the quality of the model (in terms of goodness of fit and confidence in EUI predictions).

To estimate the expected change in EUI due to a change in building technology, we first fit the model using the buildings in the user-selected peer group. The algorithm computes various coefficients associated with each of the predictors in the model. We use these coefficients to predict the baseline EUI by treating each building in the peer group as though it had the baseline equipment. Similarly, we predict the "upgraded" EUI by treating each building as though it had the post-change equipment. We compute the expected savings by subtracting the upgraded EUI from the baseline EUI, and normalize it by the baseline EUI. This results in a histogram representing the probability distribution of EUI changes due to changing the specified equipment, or the likelihood of achieving different levels of EUI changes.

To illustrate the savings estimation algorithm, consider an owner of a portfolio of big box retail buildings that is interested in investing in an equipment retrofit. The owner utilizes the BPD website to construct a peer group of buildings with the facility type "Retail – Big Box (> 50K sf)", then chooses to view savings predictions for changing heating type from "Furnace" to "PTAC" (Packaged Terminal Air Conditioner). Figure 3 shows the resulting EUI reduction histogram for the peer group of 1,186 buildings. The horizontal axis shows the EUI change, normalized by the baseline EUI and expressed as a percentage. The vertical axis is the number of buildings. The quartiles shown in the histogram indicate that three quarters of the buildings reduce EUI more than 7%, half of the buildings by more than 10%, and one quarter of the buildings by more than 13%. Similarly, the stakeholder could use the BPD app to find that 92% of the buildings reduce EUI more than 5%, or that only 2% of buildings reduce by 15% or more.

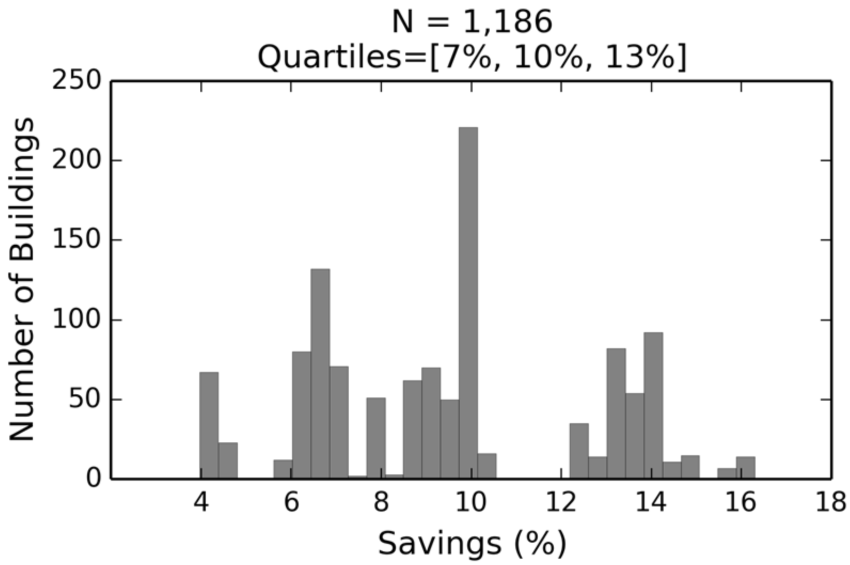

Figure 3: Histogram of EUI savings for big-box retail buildings when retrofitting heating type from Furnace to PTAC. Sample graphic; BPD user interface is being updated to include this analysis method.

The percentage of buildings in the peer group that achieve a particular level of consumption reduction can be interpreted as the probability of achieving that level of savings. The histogram can therefore be used directly to answer several questions about the expected savings, such as: "What is the probability I will reduce energy use by at least 15%?" or "What level of savings do I have a 90% likelihood of achieving?" The answers to these questions are crucial for analyzing the risk involved in energy efficiency investments. For example, if a stakeholder knows the probability that energy reductions will be greater than some threshold and can convert those energy savings to a financial measure, then the stakeholder can answer questions such as "If I invest $10,000 dollars on retrofitting the lighting in my building, how much profit can I expect?".

Testing for Representativeness

Because the BPD acquires data through voluntary submission of data collected by other entities, the database will likely not be representative for all possible peer groups. However, as the size of a peer group increases, it becomes more likely that the sample will represent the underlying population. LBNL is currently exploring ways to test how well any given BPD peer group represents the underlying building stock. We discuss two methods including: 1) comparing summary statistics against CBECS and RECS, and 2) observing the stability of the sample. The first method applies to peer groups defined at the national or census region levels, while the second applies to more narrowly defined peer groups.

The first approach compares summary statistics of a BPD peer group to CBECS and RECS. CBECS and RECS are statistically representative at the national and regional levels and are not necessarily representative for local or narrowly defined peer groups. If we choose national or regional peer groups from BPD, we can compare them to CBECS and RECS. For example, if we consider EUI for office buildings nationwide, the BPD peer group contains more than 11,000 buildings. which is 1.4% of the entire population of office buildings in the US.

Comparing these figures against CBECS, we find that the BPD's office buildings have a median EUI equal to 68 kBtu/ft^2-year, which is a 6.7% difference from CBECS median energy use intensity for office buildings. These EUIs are close, and with further analysis of subpopulations within the peer group, we may find that the BPD contains a representative sample.

Within the broad peer group "office buildings", the BPD contains data-rich subsets of buildings, for example office buildings larger than 500,000 ft^2. Shown in Figure **4**, the BPD peer group consists of over 1,000 buildings out of the estimated 8,000 buildings in the stock represented by CBECS nationally, accounting for 14% of the underlying building stock. CBECS collects only 72 samples to represent all 8,000 buildings for its national summary. Because the BPD contains a larger sample size with energy use data collected as recently as 2013, it is likely to be a better source of information about these buildings.

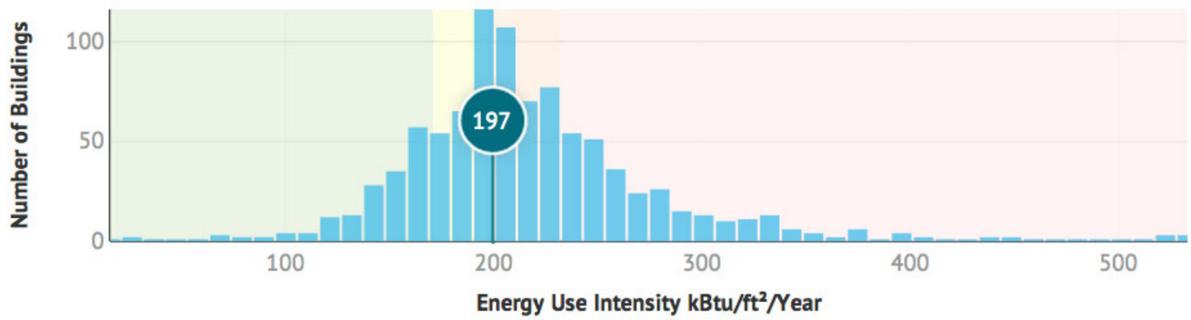

Figure 4: Histogram of energy use intensities for 1,089 office buildings larger than 500,000 square feet nationwide, with median energy use intensity equal to 197 kBtu/ft^2-year. (Source: Building Energy, Inc.)

In order to determine whether a peer group sample is representative, we begin by observing how stable the mean and variance are for the peer group. In other words, do the mean and variance change dramatically if we limit the peer group to fewer buildings? We observe this by computing the mean and variance for many random samples of buildings taken from a peer group and computing the change in the mean and variance as a function of the number of buildings in the sample. This procedure is commonly referred to as bootstrap sampling with replacement (see for example Wang, et. al. (2014)).

As the size of the sample approaches the size of a population, variation in summary statistics will likely decrease. For example, if a population contains 1,000 buildings, then summary statistics will likely not change significantly if we increase the number of buildings sampled from 998 to 999. However, if for the same population we increase the sample size from 1 to 2 buildings, then we might expect to see more variation in summary statistics between the two sample sizes, indicating that the sample is not stable with only 1 or 2 buildings. From these results, we can conclude that the samples with only 1 or 2 buildings are not representative of the population, but that the samples with 998 or 999 buildings likely are representative. When we find that the mean and variance of a small sample of buildings do not differ from the peer group mean and variance, we may deduce that the number of buildings in a peer group is adequate to represent the underlying building stock. We are in the process of confirming the degree to which this conclusion is correct. Ultimately, we may be able to use this technique to extrapolate to how many samples are needed to represent a particular group of building stock. The results of this analysis could answer questions such as: What level of compliance with energy disclosure ordinances (e.g. 100% or 70%) is needed to obtain a representative sample of the building stock? Conducting the analysis to answer these questions will be the topic of future research efforts.

Discussion and concluding remarks

Over the next decade, the BPD aims to become the authoritative resource for understanding trends in building performance. The project's goal is to achieve tens of thousands of direct users and hundreds of API users. In order to provide value to these users, the BPD has set a target of containing information on 20% of the commercial market, or about 1 million buildings (EIA 2003). The residential market is larger but less diverse and is seeing a slower pace of data collection and sharing, so the goal is 5% of the residential market, or nearly 6 million buildings (EIA 2009).

A significant barrier to realizing the full potential value of the BPD is the quantity and detail of the building records it contains. Insufficient or sparse data is a potential source of uncertainty in defining peer groups and calculating energy reductions.

EPA Portfolio Manager has achieved significant market penetration, as evidenced by the number of building owners tracking information such as building age, size, location, use and energy consumption. These data constitute minimum requirements for inclusion in the BPD, and nearly 70% of database records include little more detail than the fields listed here. These data allow BPD users to understand whether their building is a high or low performer compared to local peers. While this is an essential first step, more information about buildings' physical characteristics is needed in order to understand what is driving that performance, and identify opportunities for capital improvements. More granular data will create new opportunities for statistical analysis of the data, and allow users to control for differences such as operating hours, occupancy and building assets when evaluating building performance.

The next step in developing performance-based approaches to energy efficiency is to drive greater data collection at scale, with a specific focus on asset information as well as tracking all the information over time. Longitudinal building performance data will make it possible for the BPD to analyze the same records "pre" and "post" capital and operational changes, while controlling for other factors.

Numerous recent technology, policy and market drivers promote widespread building data collection. Technology drivers include increased prevalence of interval meters and software for building management and auditing, which make data collection the norm in new and retro commissioned buildings. Policy drivers, including building performance disclosure laws and energy efficiency incentive programs, now require data collection for compliance or participation. Additionally, utilities and local governments are eager to combine the building data they collect to facilitate better design and deployment of energy efficiency programs and policies. Finally, major companies, industry organizations and non-profits such as the American Institute of Architects (AIA), the Building Owners and Managers Association (BOMA), the US Green Building Council (USGBC) and other market players have recognized the need to track information about their member's buildings, and initiated efforts to do so.

The BPD is also part of a greater DOE effort to build foundational tools that facilitate the growth of data-driven methods of evaluating building performance, including the BEDES standardized data format, as well as commercial and residential asset scoring tools. If BEDES and the asset scoring tools gain traction in the market, more data will become available, increasing the pool of potential contributors of data to the BPD. The DOE is also committed to making data collected through its programs available through the BPD, including data from the Commercial Building Energy Consumption Survey (CBECS) and the Residential Energy

Consumption Survey (RECS), the Better Buildings Challenge, the State Energy Program, and other DOE-funded research and grants.

But motivating other entities to volunteer data for the BPD continues to be a challenge. Most data contributors share their information because they believe in the public policy goal of making empirical information about energy performance available to the public and market at large. As a result, the BPD benefits from the network effect—the more entities participate, the greater its value, which in turn encourages more participation. As extra incentives, data contributors receive a cleansed, validated and reformatted dataset with some basic analysis from LBNL, and discounted access to the commercial-scale API when it becomes available.

While the BPD is constantly seeking new data contributors, we seek to maintain a balance between data quantity and data quality. We could, for example, make it possible for users to directly enter data into BPD, but their data could be simply estimated, or using design specifications or projections. Instead, we take datasets only if we know all the buildings are real and the data collector confirms a genuine effort to enter accurate information.

The BPD will always provide a basic public user interface, but ultimately we aim to support the growth of a community of API users that integrate BPD data into their own research and products, and combine it with other public and private datasets. Planned additions to the public interface include: unlocking new filters as new fields become more common in the underlying dataset, advanced statistical analysis tools, and addition of financial assumptions in retrofit app.

API functionality was released in Spring of 2014, and was featured in the DOE "Apps for Energy" Contest, where teams of developers combined the BPD with other publicly available APIs, to develop applications for the web or mobile phones. Examples include an application that estimates energy cost of a given real estate listing using the BPD and Zillow, and an app that converts Green Button data into a "how to save" report using a BPD comparison dataset. API users could also integrate the BPD with privately held data in order to show their own buildings compared against BPD peers.

If performance-based approaches reach scale in the market, it will be easier to track information throughout a building's life cycle—from construction (modeling and code compliance) through operation, renovations, leasing and sales. In addition, this approach will achieve better integration of datasets with different levels of data granularity, from whole building benchmarking, to equipment performance and controls and transactions with the grid. Widespread collection of building performance data will make the BPD an increasingly useful tool, facilitating development of a range of other tools and services to conduct analysis at higher accuracy and lower cost.

Acknowledgements

This research was supported in part by the Building Technologies Office, Assistant Secretary for Energy Efficiency and Renewable Energy of the U.S. Department of Energy, and performed under U.S. DOE Contract No. DE-AC02-05CH11231.

References

DOE (Department of Energy). 2014a. Buildings Performance Database (BPD). Accessed May 16, 2014. http://energy.gov/eere/buildings/buildings-performance-database

DOE (Department of Energy). 2014b. Building Energy Data Exchange Specification (BEDES). Accessed May 16, 2014. http://energy.gov/eere/buildings/building-energy-data-exchange-specification-bedes

EPA (Environmental Protection Agency). 2012. ENERGY STAR Portfolio Manager Data Trends: Benchmarking and Energy Savings. Accessed February 27, 2014. http://www.energystar.gov/buildings/sites/default/uploads/tools/DataTrends_Savings_20121002.pdf

EIA (Energy Information Administration), 2003. Commercial Building Energy Consumption Survey 2003. Accessed February 27, 2014. http://www.eia.gov/consumption/commercial/

EIA (Energy Information Administration), 2009. Residential Energy Consumption Survey 2009. Accessed February 27, 2014. http://www.eia.gov/consumption/residential/

Wang, Y., M.D. Sohn, Y. Wang, K.M. Lask, T.W. Kirchstetter, A.J. Gadgil. 2014. "How many replicate tests are needed to test cookstove performance and emissions? – Three is not always adequate." *Energy for Sustainable Development* 20; 21-29. http://dx.doi.org/10.1016/j.esd.2014.02.002