MEMORANDUM

To: NEEA & the RBSA II Sampling Working Group

From: Bonneville Power Administration Market Research Team

Date: March 8, 2019

Subject: RBSA II Re-weighting Analysis & Results

Introduction

This document provides the analysis and results from The Bonneville Power Administration's (BPA's) assessment of the Second Residential Building Stock Assessment (RBSA II) sampling and weighting approach and presents new weights that can be used with the RBSA II database to calculate regional housing stock characteristics. Cadeo (the research team) supported BPA in this analysis. This work was borne out of an initial assessment of similarities and differences between the Northwest Energy Efficiency Alliance's (NEEA's) 2011 and 2016 Residential Building Stock Assessments (RBSAs, also referred to as RBSA I and RBSA II, respectively). That original comparison identified the evidence of unintentional bias in the RBSA II sample, which can affect the representativeness of the results of the RBSA study. The goals of this work were to determine the impact of the observed sample bias and propose a re-weighting approach to address the issue.

This memo describes the analysis and results of the re-weighting approach. Illustrations and results focus on the single-family study for simplicity and to provide specific examples to highlight analysis decisions. The team completed similar analyses for the manufactured and multi-family studies. Details on any differences in approach for manufactured housing and multifamily are provided in Appendix A.

Background

In reviewing the RBSA II data, the research team observed evidence of an unintentional bias in the RBSA sample that was not corrected for in the current post-stratification weighting scheme. In general, RBSA II's sampling approach divided the region into seven distinct sub-regions, each with unique sampling targets: Puget Sound, Western Washington (excluding Puget Sound), Eastern Washington, Eastern Oregon, Western Oregon, Idaho, and Montana. RBSA II also included small, utility-funded oversamples, which were essentially treated as separate samples, drawn independently from the NEEA-funded Core sample. The complete sample was then merged together at the end into the final realized sample, which was subdivided into 21 discrete strata, which are a combination of the seven original NEEA subregions, divided into utility-specific portions to account for the utility oversample.[1] The RBSA II contractor originally post-

[1] The individual NEEA subregions are divided into BPA and non-BPA portions, except Washington's Puget Sound region which is subdivided into PSE, Snohomish PUD territories, BPA and several Seattle City Light strata.

stratified and assigned weights to the sampled sites based on the inverse of the realized selection probability[2] within each of these strata.[3] However, these sample strata regions are still relatively large, heterogeneous areas. We observed that sampled sites were concentrated in specific areas within each of the seven large RBSA II sample strata (i.e., subregions), rather than equally dispersed, as would be expected with equal probability sampling. Under independent equal-probability sampling each county (and each utility service territory) within a given subregion will, on average, get a share of sample points equal to its housing unit/residential customer count, but with a lot of variation around that average.

Observed Bias

To evaluate the RBSA II sample against likely "equal probability" outcomes, the research team simulated 1,000 random trials of each sub-region's NEEA sample and looked at the distribution of outcomes compared to the actual sample distribution. In the figures below, the horizontal red bar represents the actual NEEA sample count and the blue box plot displays information about the distribution of the group's sample count across the 1,000 simulations.[4]

Figure 1 shows simulation results for the Eastern Oregon sub-region at the county level, which illustrates the trend towards higher than expected sampled sites in "large," populous counties[5] and lower than expected sampled sites in "medium" and "small" counties. Specifically, in Eastern Oregon, Deschutes County, a single county with 35% of Eastern Oregon's housing units, got 59 single-family site visits in RBSA II, but in 1,000 simulated runs it got a median of 37 sample points and never got above 53.

---

[2] The realized selection probability defines how likely each sampled site was to be selected based on the RBSA II subregional sample frame targets and how the population is distributed within each of the RBSA II subregions. Essentially, it describes the "likelihood of being picked" under a purely random, equal probability sampling approach and re-weight sites to better reflect how common they are in the region.

[3] RBSA II's sampling approach is described in the RBSA II documentation materials available on NEEA's website (https://neea.org/data/residential-building-stock-assessment), as is summarized and compared to the RBSA I approach in a previous report (Bonneville Power Administration, August 15, 2018. *RBSA Review Findings Memo.* Prepared by Cadeo on behalf of the BPA Market Research Team).

[4] In an individual box plot, the central box represents the "interquartile range", which contains half the simulation outcomes, from the 25th to the 75th percentile. The horizontal bar bisecting the box is the median of the 1,000 simulation outcomes. the vertical bars on the top and the bottom of the box, representing 1.5 interquartile ranges, can be thought of as the "two standard errors" of conventional hypothesis testing, that is, the simulated value is contained within the combine vertical bar span "almost all" of the time. The individual dots above and below this range represent the rare individual outcomes from the 1,000 simulations which lie outside this vertical bar span.

[5] To better interpret the results, we have put the counties into three categories ("small", "medium", and "large") and display the results using those groups, not by individual counties. We developed these groupings based on the population of each county. It was not possible to develop strict, quantitative thresholds to define "large," "medium," and "small" counties due to the differences in relative and absolute county populations in different subregions. Therefore, these designations were developed specifically for each subregion, where the most populous county represented the "large" county (typically only one, unless there were two or several counties with similarly large populations), the "small" counties represented the counties with the smallest populations, and "medium" represented everything in the middle. The designations are somewhat arbitrary, as they are only meant to be used for illustrative purposes in observing the trend in sampling rates between relatively populous and relatively less populous counties.

Figure 1. Country-Based Simulation Results for the Eastern Oregon Sub-Region

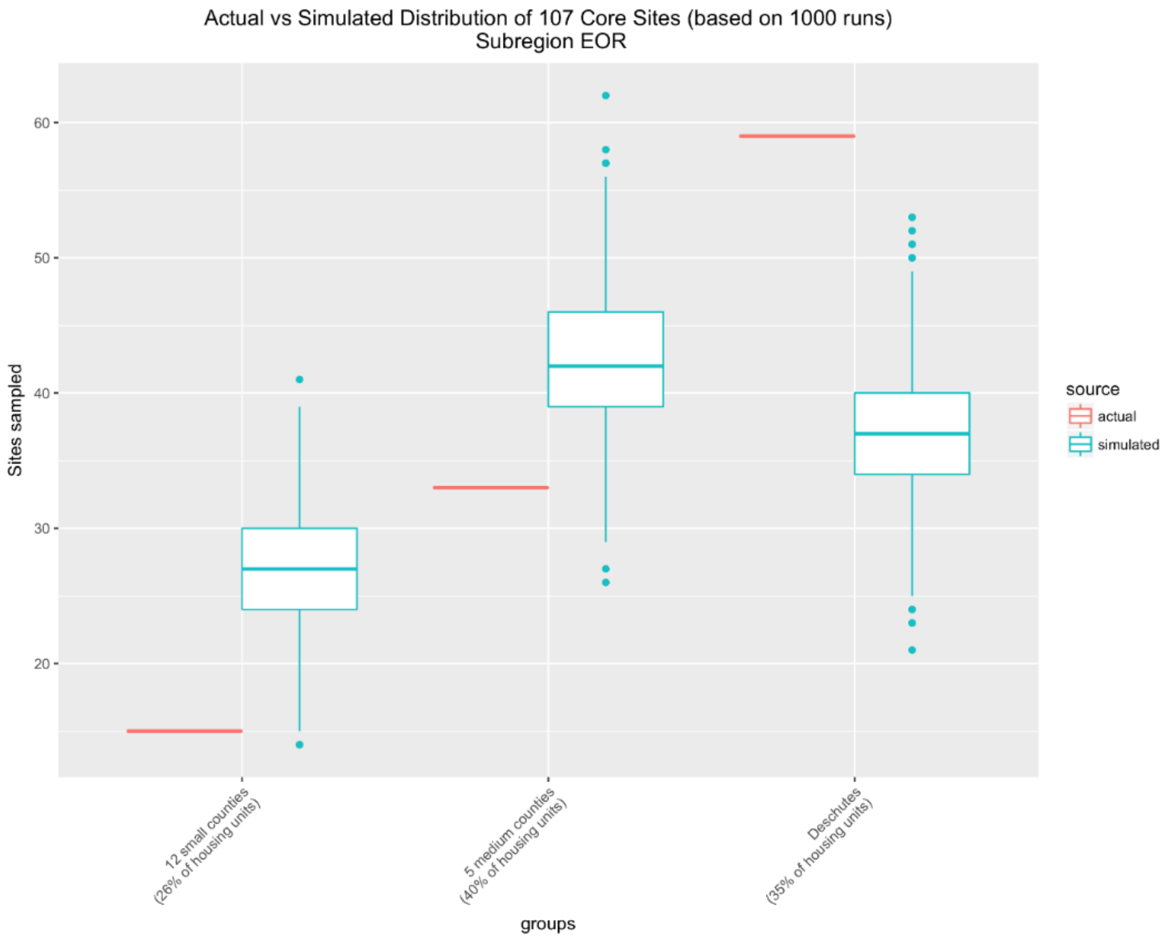

One outcome of the concentration of sample sites in large, populous regions is that utilities that serve those populous regions are similarly over-sampled and utilities that serve less populous regions are similarly under-sampled. Figure 2 displays the utility service territory version of the Eastern Oregon simulation, which displays a similar trend towards large, IOU utilities and away from smaller, public utilities. Specifically, in the chart below the biggest Eastern Oregon Utility, Pacific Power, with 45% or Eastern Oregon's residential customers, is overrepresented, with 73 single-family sites; in 1,000 simulation runs it received a median of 48, and never received more than 66. The fifteen smallest utilities, accounting for 32% of Eastern Oregon's residential customers, received 15 NEEA sample points; in the 1,000 simulations, they received a median of 34 and never received less than 18.

Figure 2. Utility Service Territory Version of the Eastern Oregon Simulation

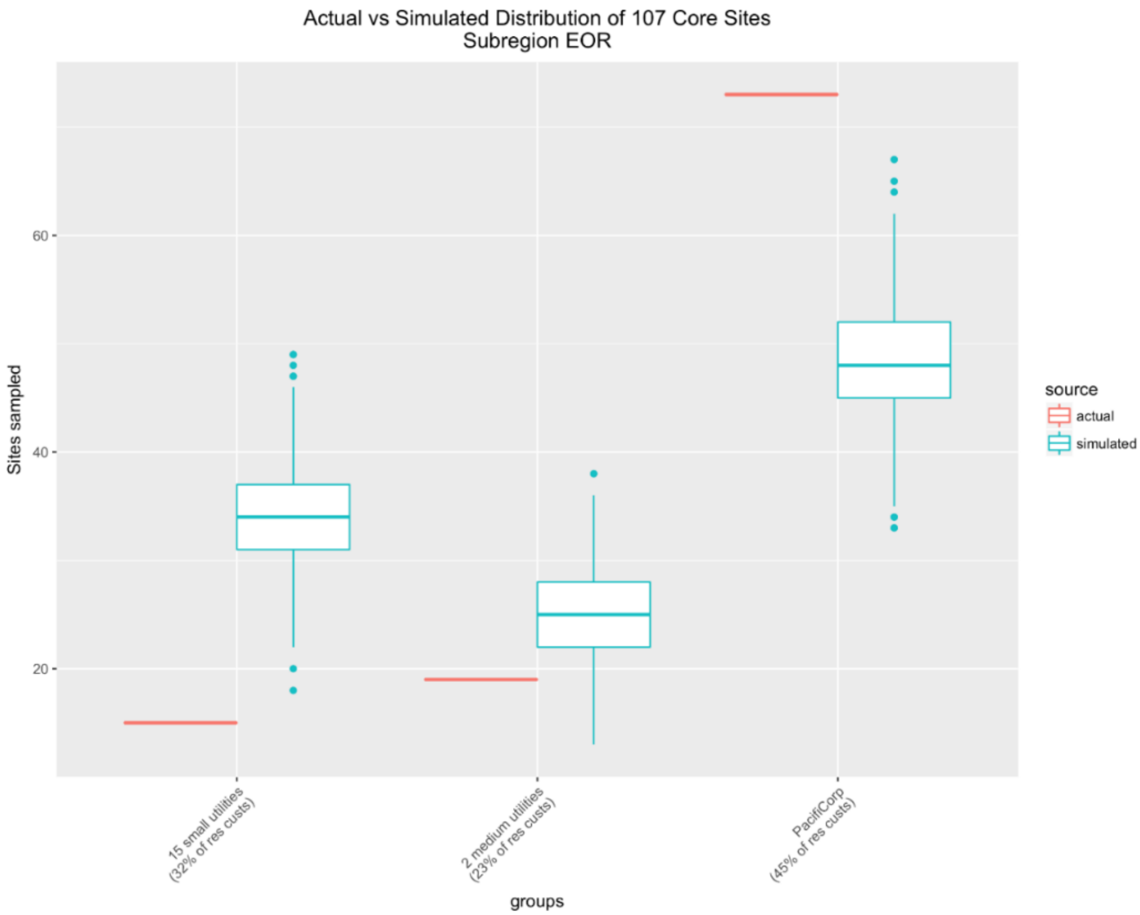

In short, these two ways of evaluating the Eastern Oregon sub-region—by county and by utility service territory—both demonstrate a departure from equal probability sampling. With equal-probability sampling, we naturally expect to find greater sample concentrations in more-populous areas. The simulation demonstrates that the RBSA II sample expressed greater concentrations than random selection would typically yield.

The other RBSA II sub-regions also display consistent geographic concentration bias, which results in a site selection preference for populous counties and for customer-rich utilities. For six out of the seven of the RBSA II-defined subregions, sites in populous counties (suggestive of urban centers) were selected over less populous (i.e., rural) counties at a significantly higher rate (significant enough to reject equal probability sampling at the 95% confidence level). The exception is Puget Sound, which still exhibits the urban/rural bias observed in other counties, although at a less striking rate,[6] potentially due to the relatively large populations that exist throughout the county. Montana also illustrates an interesting pattern of selection for "medium" counties, at the expense of both smaller and larger counties. This trend is demonstrative of the

[6] Puget Sound exhibits a mild preference for King county, the most populous county; a median sample rate for the three other populous Puget Sound counties (Snohomish, Pierce, Thurston) and a disproportionately small number of sites for the relatively rural Island, San Juan, Skagit, and Whatcom counties.

geographic concentration bias and departure from equal probability sampling, but reflects a preference for, potentially, *certain cities* in Montana, as opposed to cities in general.

As discussed earlier, concentration of sample sites in populous areas also results in concentration of sample in utility service territories that serve those populous areas. Because IOUs typically serve large, urban population centers, for five of the seven RBSA II subregions, IOU sites were selected at a significantly higher rate than public sites (significant enough to reject equal probability sampling at the 95% confidence level). The exceptions are Puget Sound and Western Washington subregions,[7] both of which have large public utilities (Clark PUD and Seattle City Light[8]) that serve the most populous cities in those subregions.

The implication of this finding is that the current RBSA II weights over-represent the sites from heavily-sampled areas and under-represent the sites from less-heavily-sampled areas. This bias could impact the representativeness of the study results for the region, as any differences in building stock characteristics between the more and less heavily sampled areas (due to, for example, differing climates, rate environments, fuel availability, or vintage construction practices) would be skewed toward the more heavily sampled populations.

The reason for the observed bias is unknown, although it could be influenced by the source of the sample frame (the postal service addresses), the method of selecting sites from the sample frame, the method of recruitment, and/or the method of replacement for refusals. The goal of any sampling methodology is to control for selection bias in the sample execution (i.e., selection and recruitment). However, this can be challenging to do due to the realities of self-selection and low response rates;[9] post-stratification can also be used to identify and account for any apparent bias in the realized sample, which is what we propose to do here. Because we do not know the root cause of the underlying bias, we do not know for certain that population density caused the bias to occur, we can only observe that the bias is correlated with population density. It is possible that some other factor, correlated with population density, is affecting the sample that we are not correcting for. By adjusting for the observed bias based on population density, we are better aligning the realized sample with the likelihood of sample selection based on equal probability sampling. However, we do not know if there may be some *other* characteristic, such as socioeconomic status or demographics, that is actually driving the sample bias, and which is not fully accounted for by the proposed geographic concentration post-stratification method.[10]

Methodology for Post Stratification and Recalculated Weights

To account for the observed geographic concentration bias, the research team:

---

7 As described in the RBSA II report, the RBSA II sample consists of seven unique subregions: Puget Sound, Western Washington (excluding Puget Sound), Eastern Washington, Eastern Oregon, Western Oregon, Idaho, and Montana.

8 Note Seattle City Light received an oversample, but this analysis does not account for those additional sites.

9 There are many potential sources of bias in large-scale field studies that can affect the accuracy and comparability of the results. For example, low recruitment or only recruiting certain segments of the population could affect the representativeness of the results.

10 RBSA II collected more detailed demographic information, but did not use these data to assess the representativeness of their sample population from a demographic perspective. Future studies should consider the impact of response bias and significant shifts in demographic information in the comparability and representativeness of the results. Specifically, additional investigation into potential response bias based on demographic comparisons to Census data, additional data collection from publicly available sources (such as assessor data) on non-responsive homes, and other bias studies would be extremely helpful to help understand if and how the sampled set of homes is similar to and representative of the non-sampled homes and sample strata as a whole.

- Characterized the data sources and developed a quantitative scale for population density and approach for reweighting the sample at a more granular level
- Post-stratified the RBSA II sample using the new post-stratification approach (discussed in detail later in this memo)
- Recalculated results for select key variables using new weights and original RBSA II weights to determine the impact of the bias and new alternate weights on the RBSA II results

The new post-stratification method resolved the geographic concentration bias by dividing up the region and the RBSA II sample into more granular sample strata within each large RBSA II subregion. The new sample bins are defined based on the county and utility type (i.e., BPA and non-BPA utilities).[11] That is, a single sample bin would comprise the BPA portions of a given county and a separate sample bin would comprise the non-BPA portions.

The new sample weight for each of these new, more granular strata was then calculated based on the inverse of the realized selection probability for each bin. That is, the new sample weight for each sample bin was calculated as the total population of a given sample bin divided by the number of sample sites in that bin and can usually be thought of as the number of population sites 'represented' by a given sample site. Separate sample weights were calculated for each building type (single-family, manufactured, or multifamily). The sample weight for each bin was calculated based on the building-type specific household count based on zip-code level data from the 2016 U.S. Census, the American Community Survey (ACS) and a zip code utility service territory correspondence file. Specifically, each county and utility type-defined sample bin comprises one or several zip codes. For each zip code, the household count was calculated based on the total household count from the 2016 Census, scaled for each specific building type (e.g. single-family, manufactured home, multifamily) based on data from the American Communities Survey, as shown in Equation 1.

Equation 1: Post-Stratification Household Count Calculation

$$HousingCount_{h,s} = \sum_{c=1}^{n} CensusHousingCount_c * \frac{ACS\ HousingCount_{h,c}}{ACS\ HousingCount_c}$$

Where:

h = building type, e.g., single-family

s = sample bin (defined by county and utility type)

c = zip code

n = total zip codes in each sample bin (s)

Equation 1 also shows how the zip code-level household counts were aggregated into county and utility-type defined bins based on a zip code correspondence file which maps particular zip codes to utilities and counties. The research team used a modified version of the RBSA I zip code

[11] The team had originally intended to also include an "urban" vs "rural" indicator explicitly in the post-stratification approach, but this was not possible because of the availability of data to describe this variable for each sample bin (urban/rural designations is only available from the Census Department at the county or census block group level, not service territory or zip-code level, which is necessary for this analysis) and the fact that the sample bins would get too small to analyze, especially given the correlation and interdependent nature of the urban/rural variable with the other two variables: county and utility type.

correspondence file because certain attributes of the RBSA II zip code correspondence file made it more difficult to use for this analysis.[12] In some cases assignment to a single county was made based on population concentrations within a zip code area where the zip code did not fall cleanly in one county.

The calculated sample bins were not used directly as strata because nearly half of them (101 of 223 for single-family) did not contain sampled units and another 33 had a single sampled site.[13] Instead, the team grouped individual sample bins together based on their realized sample rate (i.e., how heavily or lightly sampled the bin was in the actual RBSA II sample). The team also accounted for geography (to account for any weather- or housing vintage-related dependencies), utility type (to account for the availability and penetration of heating fuel), and population (a proxy for urban/rural or population density) in aggregating sample bins into new sample strata for analysis. This approach sought to group counties with similar geographic, utility service territory, and population density characteristics together if they were not sampled at greatly disparate rates. This approach achieves the goal of fairly homogenous strata with similar sampling rates, utility types, and populations, which we believe should result in results for each new sample strata that are more representative of the sample population due to the more homogenous characteristics of each strata.

Validating the Re-Weighting Approach

Because the approach for aggregating sample bins into strata described above (Approach 1) requires some judgment, the team tested the sensitivity of the results to the selected aggregation method to ensure such judgment did not have a disproportionate effect on the overall outcome of the post-stratification. The team developed a second approach (Approach 2) to validate Approach 1. In the second approach, the team separately sorted county-level BPA and non-BPA sample bins by single-family household count and then aggregated those into larger bins based on the single-family housing count alone, irrespective of sampling incidence or geographic proximity. This resulted in sample strata that had similar utility type and similar population, but does not account for adjacency of counties and does not address sampling incidence directly. Approach 2 is more mechanical and objective than Approach 1 but it still has components that require judgment, such as a predefined number of strata and population cut-off points.

The benefit of Approach 2 is that it creates a comparison for Approach 1 to validate any judgment used did not disproportionally impact the post-stratification results.[14] In the next section, we provide comparative RBSA analysis results. Results for Approach 1 and Approach 2 are similar, which demonstrates that the major impact on the sample weights is derived from the more granular and detailed nature of the new sample strata, which results in a more consistent

[12] The RBSA II zip code correspondence file contained many zip codes that were mapped to larger entities than the individual serving utility, such as "Bonneville Power Administration" or "Energy Trust of Oregon." As this analysis required mapping to specific utility so that each one could be treated separately, the RBSA I correspondence file was updated. The slight differences between these files are not expected to significantly affect the analysis, as the majority of households are still mapped consistently in both files. However, for the zip codes that are not mapped consistently, the research team notes that the updated RBSA I file agrees more consistently with the zip code map available from the Regional Technical Forum (see https://rtf.nwcouncil.org/work-products/supporting-documents/climate-zones), than the RBSA II file.

[13] Extreme case: IOU-served Skagit Co, Washington with an estimated 44,000 SF housing units has no sample.

[14] Note, non-BPA utility oversample service territories were exempted from the aggregation approaches, due to their small and concentrated nature. For example, Seattle City Light comprises four separate strata and a single stratum contains Snohomish County and the public part of Island County.

and homogenous sample in each stratum than the original stratification. The strata definitions are also similar between Approach 1 and Approach 2. An example is outlined in Table 1, showing the 11 counties that comprise the single western Washington subregion in the original stratification. Those counties are divided into four (4) strata using Approach 1 and five (5) strata using Approach 2 and are, in general, categorized similarly. **In the end, the research team, in consultation with other regional subject matter experts, selected Approach 1 for the final post-stratification due to the fact that it more directly accounted for and corrected for the observed geographic concentration issue.**

Table 1. Example Strata Differences and Similarities – Original WWA Subregion

County	Utility Type	Single-family Housing Counts	Site Count	New Sample Strata (Approach 1)	Comparison Sample Strata (Approach 2)
Clark	BPA	143,709	61	**WWA.1**	WWA.1
Skamania	BPA	4,139	2	**WWA.1**	WWA.4
Clallam	BPA	28,925	10	**WWA.2**	WWA.3
Cowlitz	BPA	34,093	9	**WWA.2**	WWA.2
Mason	BPA	26,266	7	**WWA.2**	WWA.4
Kitsap	non-BPA	87,805	19	**WWA.3**	WWA.5
Grays Harbor	BPA	28,384	4	**WWA.4**	WWA.4
Lewis	BPA	25,310	4	**WWA.4**	WWA.4
Jefferson	BPA	1,116	1	**WWA.4**	WWA.4
Pacific	BPA	11,800	0	**WWA.4**	WWA.4
Wahkiakum	BPA	1,684	0	**WWA.4**	WWA.4

Results

Once the new sample strata were defined and the team had calculated the corresponding sampling weights based on the inverse of the household count for each building type, the team calculated select RBSA results using the updated population weights for comparison with the original RBSA II weighting scheme.

Applied to primary heating and water heating fuel, the new weights result in small (1-2%) but persistent shifts towards electricity, wood (home heating only), and propane, and away from natural gas, as shown in Figure 3 and Figure 4, respectively. For clarity, Figure 3 and Figure 4 focus on electricity and natural gas. The results are also shown in a table (Tab1e 2, page 10), which shows a summary of changes in primary heating and water heating fuel for electricity and natural gas in single-family homes.

Figure 3. Regional Single-Family Primary Heating Fuel Choice

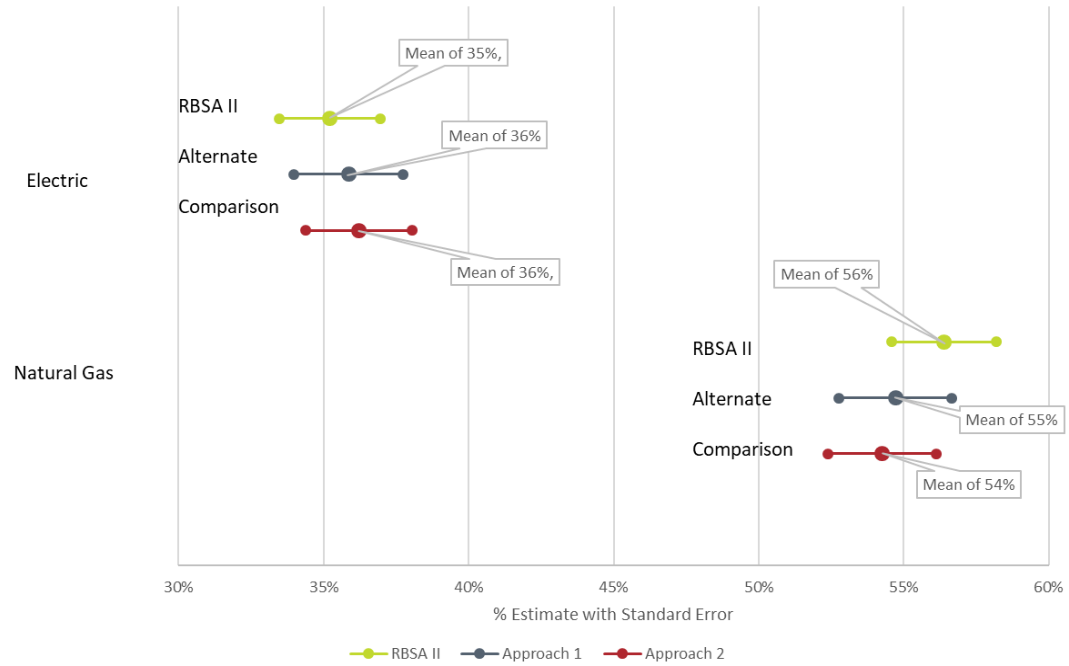

Figure 4. Regional Single-family Water Heating Fuel Choice

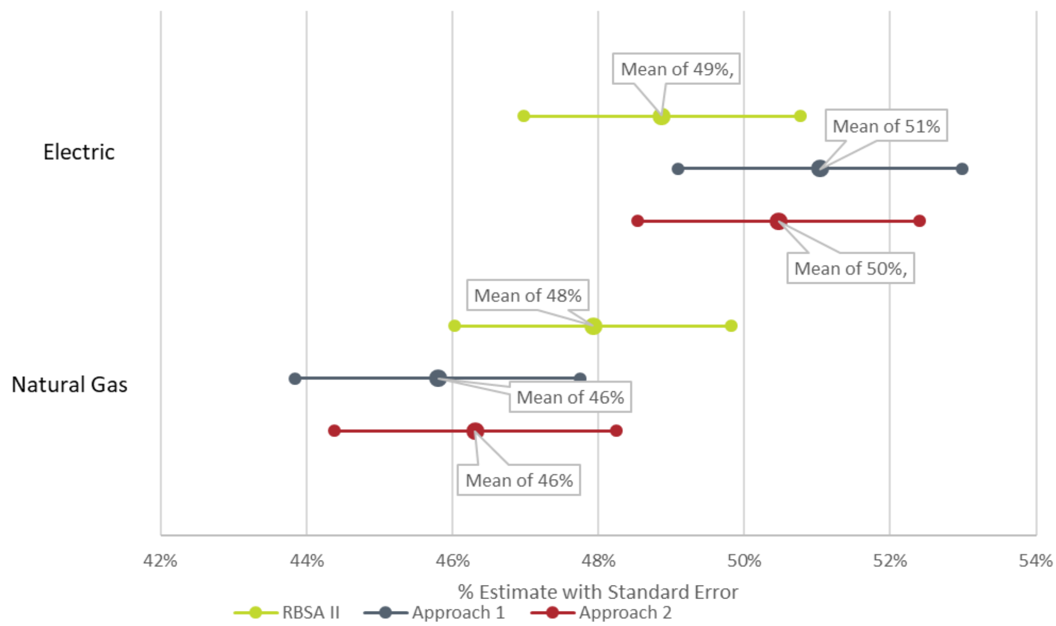

The shifts are greater at the subregional level (5-6%) and greater for some regions than others, as highlighted by the regional and western Washington results for water heating fuel shown in Figure 5.

Figure 5. Western Washington Single-family Water Heating Fuel Choice

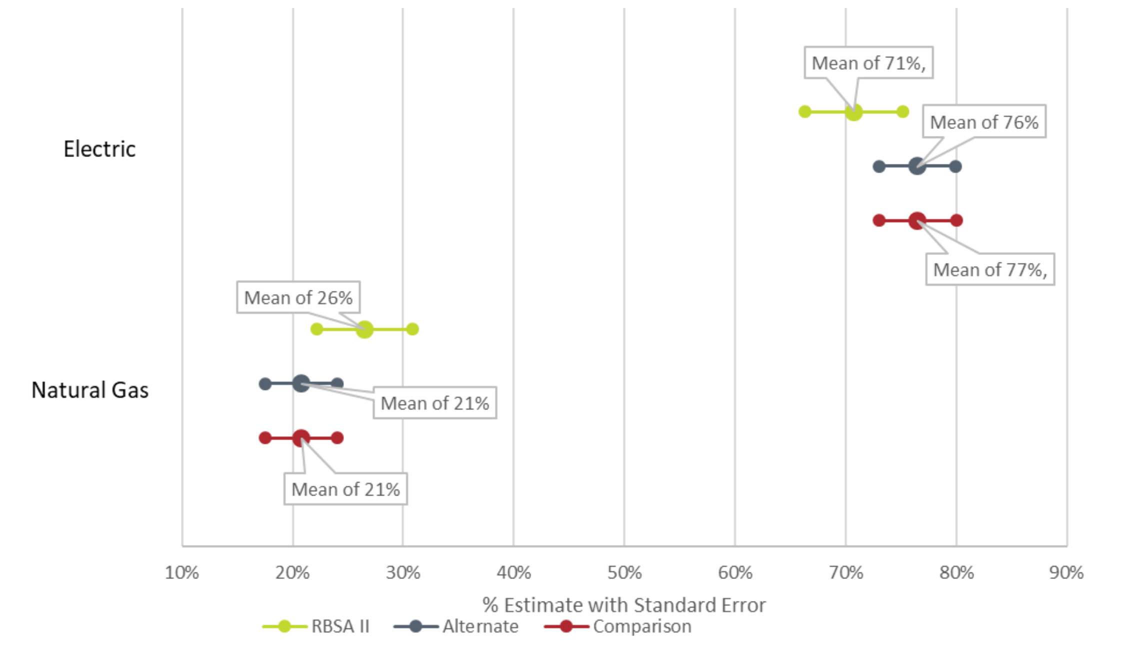

In addition, the shifts are greater for some technologies than for others, as illustrated in Figure 6, which shows that lighting results shift only slightly, if at all using the new sampling weights.[15]

Figure 6. Regional Single-family Lighting Technologies

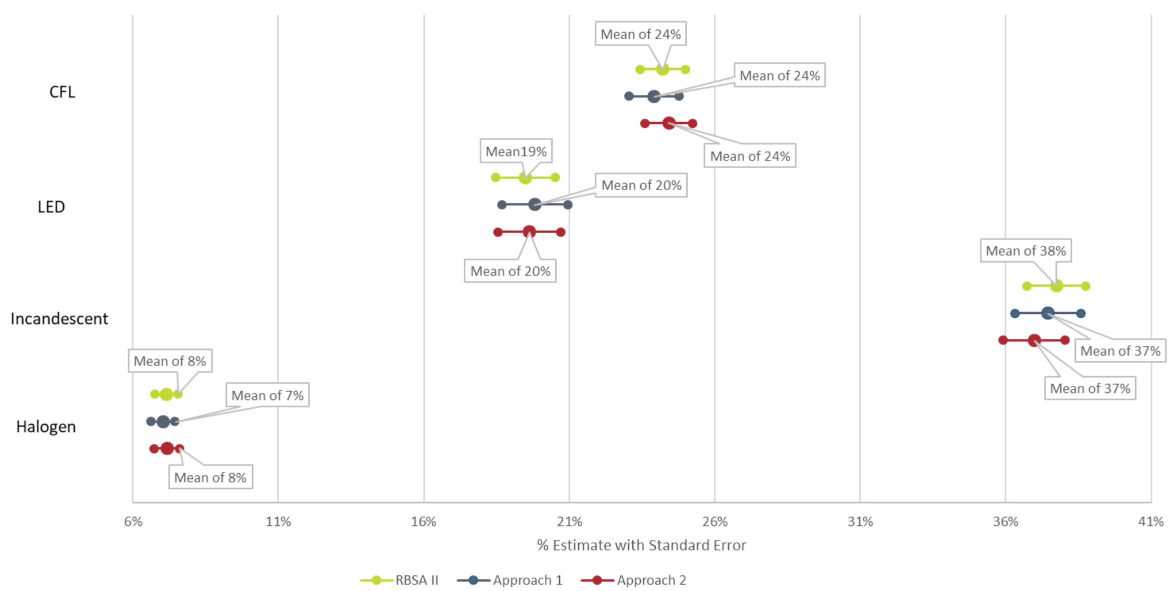

[15] We only evaluated the impacts of the new post-stratification approach on the RBSA results for three key characteristics: primary HVAC heating fuel, water heating fuel, and lighting in this analysis. However, the full set of new sample weights are included as Attachment A, which analysts can use to evaluate the impacts for all other variables.

Table 2. Summary of Changes in Primary Heating and Water Heating Fuel for Electricity and Natural Gas in Single-family Homes

Characteristics	Fuel	Original %	Original EB	Updated %	Updated EB	Change
Regional Primary Heating Fuel	**Electric**	35.2%	1.7%	35.9%	1.9%	0.70%
	Natural Gas	56.4%	1.8%	54.7%	1.9%	-1.70%
WWA Primary Heating Fuel	**Electric**	65.0%	4.6%	71.0%	3.8%	6.00%
	Natural Gas	27.0%	4.3%	22.0%	3.4%	-5.00%
Regional Primary Water Heating Fuel	**Electric**	48.9%	1.9%	51.0%	2.0%	2.10%
	Natural Gas	47.9%	1.9%	45.8%	2.0%	-2.10%
Regional Lighting	**CFL**	24.2%	0.8%	23.9%	0.9%	-0.30%
	Halogen	7.2%	0.4%	7.0%	0.4%	-0.20%
	Incandes.	37.7%	1.0%	37.5%	1.1%	-0.20%
	LED	19.5%	1.0%	19.8%	1.1%	0.30%

These are examples of results for the select variables we analyzed. The full alternate weights for each site in RBSA II database for SF, MH, and MF are provided in an Excel workbook as Attachment A. Analysts can use these to calculate or re-calculate results at a regional or subregional level and, we believe, will result in more representative values than the original RBSA II weights.

While not statistically significant from the original RBSA II weights for the variables we analyzed, we believe the changes in HVAC heating fuel and water heating fuel are large enough to warrant updating the RBSA II database to account for the new weights, as the new weights better tie the sample to the population at a more granular level. The impact of using the new weights can have systemic and long-lasting impacts on regional assessments until the next RBSA is conducted. For example, the current weights potentially understate the number of electrically-heated households in the region by 30,000 to 45,000 units, which represent approximately 60 to 65 aMW of electric load.[16]

As discussed above, this assessment does not address the root cause of the observed sample bias or how to prevent such issues in the future. A previous analysis[17] describes, in detail, the sample methodology employed by RBSA II and outlines several next steps to prevent or address such issues in future stock assessments.

[16] Calculated based on an average net change in electric load of 7,500-9,600 kWh/yr for HVAC and 3,000 kWh/yr for water heating.

[17] Bonneville Power Administration, August 15, 2018. *RBSA Review Findings Memo.* Prepared by Cadeo on behalf of the BPA Market Research Team.

Appendix A: Manufactured Housing and Multifamily Analysis

Manufactured Housing

The manufactured housing analysis and post-stratification was done identically to the single-family analysis described above. The main differences are that:

- For manufactured housing, unlike single-family, the stated sample design and size was state-wide, for both the NEEA Core sample and BPA oversample. As such, the analysis does not consider "subregions," as in the primary single-family analysis. Instead, the simulations and sample weights are redistributed across each state based on the inverse of the realized selection probability.

- The manufactured housing sample appeared to be much more concentrated in some cases and, therefore, contained many more county- and utility-type-defined sample bins that were not sampled at all. This is partly due to the significantly smaller sample, but many states have significantly more "zero count" sample bins than the equal probability simulations suggest they should. As such, the manufactured housing sample required more aggregation with only approximately 10 sample strata per state, some of which contain large areas without any sample in them. This is not ideal, but is a reality of the data.

Multifamily

The multifamily analysis and post-stratification employed a different post-stratification approach due to the differences in sample design and building characteristics for multifamily units.

Multifamily Sample Design Differences

Specifically, the multifamily analysis allocated sample weights across pairs of states (WA and OR, ID and MT) because the multifamily housing unit sample had a different sample design than either the manufactured housing or single-family sample, consisting of:

1) A NEEA Core sample of 285 units hard-allocated across states proportionally to their estimated multifamily unit population,[18] but not specified at state subregion levels.

2) A BPA oversample of 208 units hard-allocated to public subregions as follows: PS:31, WWA:46, EWA:43, WOR:47, EWA and EOR together:49, ID and MT together:35

3) A PSE oversample of 49 in Puget Sound Energy service territory.

Since the BPA oversample subregions crossed state lines in two instances (EWA-EOR and ID-MT), total state samples were not fixed beforehand, and simulations to test for equal probability sampling within each stated sample frame were actually carried out simultaneously using pairs of states (WA and OR, ID and MT). As with single-family and manufactured housing, the simulations were carried out to judge the degree to which the realized sample departed from or successfully carried out the stated sample design.

[18] (WA:172, OR:76, ID:21, MT:16)

Differences in Building Characteristics
The multifamily sample is in fact two related samples, one of multifamily buildings (common systems, public spaces, envelope), and one of individual multifamily units. The primary multifamily restratification was done at the unit level (not the building level, although those were also sampled). This analysis was done similarly to the single-family analysis described in the body of the memo, except that several buildings included multifamily units where more than one unit was sampled (3 unit triplets and 19 unit doubles).[19] This poses a weighting problem because the "extra" units (the second or third units surveyed out of sampled buildings) may not have been randomly selected in the same way as the original unit.[20] If the multiple units were recruited because they were in the same building as an already-recruited unit, the selection probability for each of these units is double or triple the background recruitment probability. In the new re-weighting scheme, we presumed that these multiple units were recruited based on proximity, due to the improbability of selecting two or three units from the same building randomly under equal probability sampling. Therefore, to account for this and ensure that the end results are not overly biased towards the sampled buildings with multiple units, the new sample weights for these units is based on half or a third of the weight that a singleton unit would get, as well as the same wider-scale departures from equal-probability selection (i.e., geographic concentration bias) seen earlier.

In addition to multifamily units, select multifamily buildings are also sampled and reported. Due to issues in availability of data and significant differences in weighting approach between units and the multifamily buildings, we did not re-weight the multifamily building sample. These issues are described below.

In practice, multifamily buildings are often recruited through the unit. This is due to the availability of recruitment data at the unit level (i.e., postal addresses or, in the case of RBSA I, phone numbers), as well as the availability of sample-frame definitional data (e.g., ACS five-year averages and annual housing unit estimates from the Census), which also provide a very direct way of estimating multifamily unit counts within chosen geographic areas. By contrast, there is no available sample frame of regional multifamily buildings and no off-the-shelf way to estimate their total count or size distribution within specified geographic areas. However, this information is necessary in order to generate sample weights at the building level. RBSA II provides estimates of multifamily building population counts for each of its strata that were calculated based on the total multifamily units (from the Census data described above) divided by the average total unit count of surveyed buildings in that stratum, assuming that is a good estimate of actual multifamily building size (unit count). However, calculating building weights this way does not account for the fact that larger (higher unit count) buildings are more likely to be selected, so the reported building-level characteristics will likely be significantly skewed towards the attributes of these larger buildings.

This results in an upward-biased estimate of average building size, and a proportionally downward-biased estimate of building population count. Precisely how large the bias is depends

[19] 47 of the 542 sampled units were in buildings with multiple units.

[20] The RBSA II documentation is not clear as to whether such sites were randomly recruited or were recruited by knocking on neighboring doors after selecting the primary unit, which is a typical recruiting practice in multifamily buildings.

on the distribution of building sizes in the population; the more spread in building sizes, the worse the bias. Some numerical experimentation suggests a bias in the range of 40% to 300%, although there is no real upper limit; it depends on the spread and skew of building sizes. One could attempt to correct for this issue by calculating *relative* probability weights for individual variables (e.g., energy use per square foot, proportions of floor area conditioned by various HVAC technologies, glazing ratios, or building size) without knowing the population count of buildings. This analysis was outside the scope of the current re-weighting scheme, but could be pursued in the future to improve the representativeness of multifamily building weights for this or future RBSAs.[21]

---

[21] Note, this relative analysis is similar to what was done in RBSA I.

Attachment A. Full Alternate Weights for All Building Types

Attached as Excel workbook, “RBSAII_new_samplewts.xlsx”

RBSAII_new_sample
wts.xlsx