B O N N E V I L L E   P O W E R   A D M I N I S T R A T I O N

Residential Ductless Mini-Split Heat Pump Retrofit Monitoring

June 2009

Prepared for:

Bonneville Power Administration
Contact: Kacie Bedney, kcbedney@bpa.gov

Prepared by:

ECOTOPE
Consulting
Research
Design

Kevin Geraghty, David Baylon, Bob Davis

Table of Contents

List of Tables

List of Figures

Executive Summary

Beginning in the summer of 2007, BPA developed an initial pilot study to provide basic information on the energy savings potential of the ductless heat pump (DHP) technologies introduced in 2006. The primary goal of this study was to explore the use of modern submetering technology to ascertain the performance of these systems in customers' homes. Specifically, the target customer group comprised of existing homes heated by zonal electric resistance heating. This initial effort included 14 sites in three utility service territories (City of Monmouth, Grant County PUD, and Tacoma Public Utilities).

The primary goals of this study were to:

- Provide an early verification of RTF energy savings assumptions.
- Gain experience for a larger review of DHP retrofit in zonal electric resistance heated homes.
- Review data collection procedures and refine the instrumentation protocol for the initial region-wide pilot project currently underway (DHP Impact and Process Evaluation).

Ecotope analyzed prior and subsequent energy use at fourteen sites that received DHP retrofit installations under this pilot program in 2007 or 2008. Two sites were in the Columbia Basin (Grant County, WA), but the remaining sites (eleven in Monmouth, OR, one in Tacoma, WA) were in western Oregon and Washington. The sites received a "quad-metering" monitoring package to record the performance of the DHP installation over the course of a year (minimum) ending March 2009. Utility bills were collected to provide the basis for estimating savings in space conditioning that resulted from this technology. The installations were designed to provide supplemental heating to offset the heating energy requirements from electric resistance zonal heating systems. The DHP was installed to displace space-conditioning requirements with a high efficiency heat pump in the central zone of each of these homes. The installation was not designed to provide a full replacement heating system.

The analysis used a "variable-base degree-day" methodology (often referred to as PRISM) to evaluate the energy requirements and the space heating usage of these homes prior to the installation of the DHP system. The quad-metering system provided a direct measure of the heating usage during the one-year study period. Results were weather adjusted to provide a valid estimate of energy savings. All savings estimates are expressed in terms of the one year of DHP operation (March 2007 to March 2008).

Our point estimate for average per-site savings in space-conditioning consumption adjusted for that year was 4,442 kWh/year. Median savings (less sensitive to possible data and estimation problems in outliers) was 2,971 kWh/year. Heavy energy use prior to the DHP installation appeared to be positively associated with realized savings. We attribute this result to a utilization effect; site savings showed little or no association with house size (square feet). This analysis is not surprising given that installed units were of the same size and not generally capable of supplying all space-conditioning demand in the recipient houses. Within the limits of available data, DHP installation did not appear to generate significant new cooling load. The results of this study suggest savings levels consistent with the original assumptions from the Regional Technical Forum (RTF). While these results should not be considered reliable to generalize the performance of this technology to the region, a larger pilot study has been undertaken to provide such an estimate. Nevertheless, this study provided very useful data and experience in developing and executing the larger regional study.

1. Introduction

This report describes the analysis of energy usage data for a small set of ductless heat pump (DHP) systems that were part of a pilot utility program in the Pacific Northwest. The data were mostly drawn from sites in Monmouth, OR (11 sites) with two additional installations in Grant County, WA and one installation in Tacoma, WA. Data loggers were installed at each site and left in place for at least one year. Consumption related to hot water heat, to resistance heat, to ductless heat pump (DHP) use, and whole house load, were recorded separately as hourly averages, as were indoor and outdoor temperatures.

The goals of this initial pilot study were to:

1. Provide a proof of concept for the "quad-metering" protocol using real time monitoring of kWh energy use and temperature. This approach was meant to expand on the "triple-meter" protocol used in BPA residential programs throughout the 1980s and early 1990s. Like these programs, space conditioning is measured separately from domestic water heating (DHW) and total consumption. Unlike that protocol, the measurements are logged in real time and are designed to separate the consumption of the DHP systems from the remaining electric resistance zonal heating system.
2. Assess the viability of using billing data collected from a period prior to the installation of the quad-metering system to provide the baseline consumption used for calculating changes in space-conditioning consumption. This approach included the use of classic weather adjustment procedures on billing data to develop weather compatible consumption estimates for purposes of estimating heating energy savings.
3. Develop changes in the metering protocol and data collection that would enhance the confidence and veracity of energy savings estimates from these metering and analysis procedures.
4. Summarize consumption characteristics from homes metered in this initial project.

Generalizing performance from a small group is problematic but the installations and metering results provide a very useful overview of the potential of these systems. Furthermore, this study was an experiment meant to develop and refine an initial approach to DHP metering and energy savings analysis. These efforts informed the methodology for the larger and more systematic evaluation study currently underway. There are some data deficiencies (no measurements to distinguish heating and cooling loads in these DHPs, no logged whole house consumption on two sites, not very precise installation dates). On the positive side, the research has provided savings estimates for this group and invaluable insights into the types of data that must be collected to distinguish between heating and cooling operation.

2. Data Overview

Table 1 presents summary breakdowns of submetered, post-installation consumption for the 14 sites. The two Moses Lake sites lack whole house consumption data, so baseload, which is calculated as a residual after space heating and water heating are subtracted from whole house usage, is also missing. A minimum of 13 months and a maximum of 18 months of submetered data were collected for the sites. Figure 1 presents the average information for the 12 sites that had measured whole house consumption data. Roughly 50% of average consumption was residual baseload. Resistance heat averaged about 12% of total consumption, while DHP usage was about 15% of the average total consumption. This suggests a substantial contribution to the space conditioning requirements of the house. Water heat was 24% of the average consumption. An easy mnemonic rule of thumb for the sample would appear to be half residual baseload, a quarter space heat, and a quarter water heat.

While the lack of detailed characteristics and occupancy information makes the comparison with previous studies difficult, the averages of the metered energy use for DHW and other energy use is consistent with work done under the RCDP program. With a larger and more complete study group, this data set will offer the opportunity to compare a modern sample of homes to the various BPA programs that provided the base consumption data used in program design and regional energy planning over the last two decades.

DHW results are very consistent with the results of the BPA residential programs from the early 1990s (see Quaid et al., 1991 for site built homes, and Roos & Baylon, 1993 for manufactured homes). For the other non-space conditioning loads, this group of homes suggests that about a 25% increase results from these previous submetered studies. The size of this sample makes any further comparison to these studies difficult to justify, but the indication is that this approach will serve to update the data developed twenty years ago to assess the changes in consumption patterns in the modern residential sector.

Table 1: Annualized Submetered Consumption

Site	Area (Sq Ft)	DHP (kWh/yr)	Resistance Heat (kWh/yr)	Water Heat (kWh/yr)	Residual Baseload (kWh/yr)
Monmouth 1	2600	2903	3873	4663	11624
Monmouth 2	1300	2334	1	4803	12495
Monmouth 3	1800	4547	0	2284	4181
Monmouth 4	1650	3715	5080	4137	10791
Monmouth 5	1560	1543	2388	4487	5314
Monmouth 6	2025	2052	5949	3697	8299
Monmouth 7	2350	2464	2545	4998	9187
Monmouth 8	2100	5345	4633	8138	10796
Monmouth 9	3100	2825	933	4648	12766
Monmouth 10	2800	4404	2806	6340	11552
Monmouth 11	1700	2007	391	7488	11444
MosesLake 1	1100	2164	1368	3250	
MosesLake 2	1250	5355	1552	3378	
Tacoma 1	1200	1930	335	2390	8125
Average	**1903**	**3006**	**2411**	**4839**	**9714**

Figure 1: Average Post-Installation Consumption Proportions

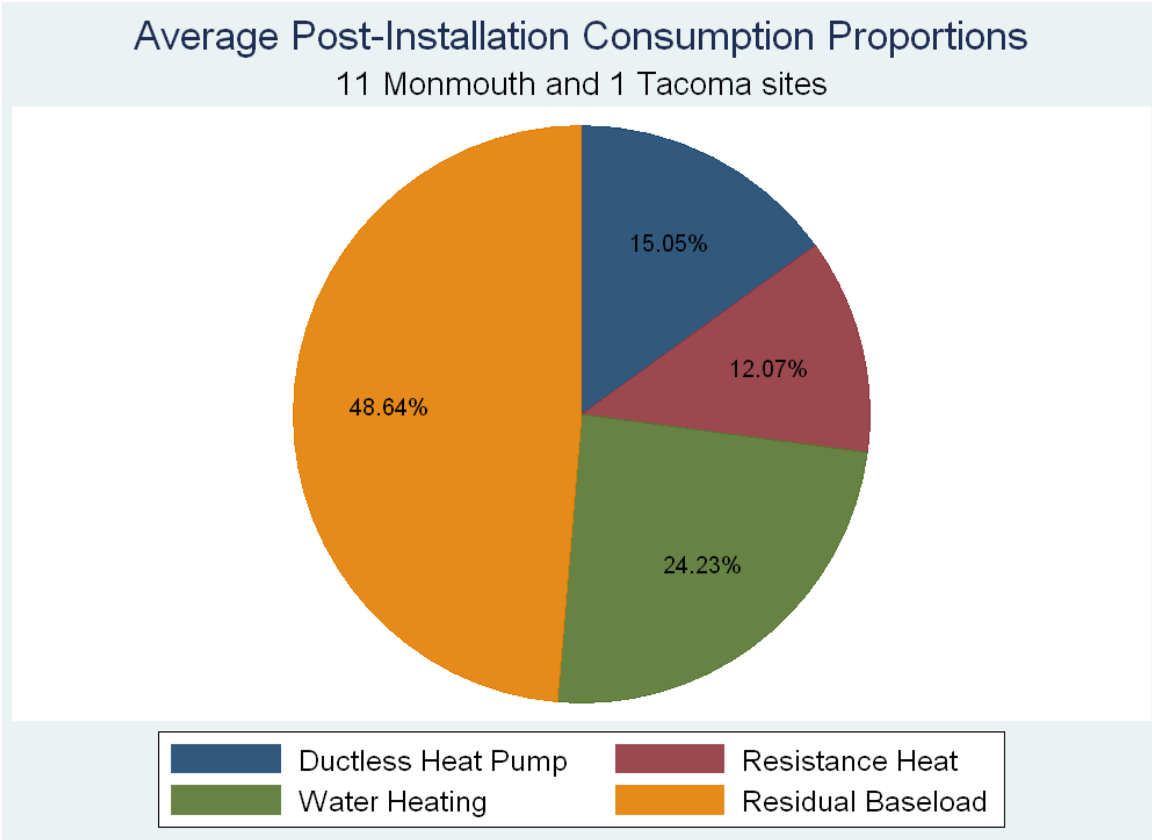

3. Savings Estimation Approach

Given the small study population, and given that the pre-installation (electric bills) and post-installation (direct submetering) data sources are disparate, the comparisons and analysis was made as simply as possible in order to keep modeling assumptions and data manipulations to a minimum. As a result, we have treated these sites as case studies, which allows us to explore all the issues that are important to the goals of this pilot project.

3.1. Energy Consumption

Up to three years of utility billing data were collected for the pre-installation period. The most recent pre-installation period was usually the period that ended in the fourth quarter of 2007. Although, in some cases only a less contemporary billing period was available. Since it is very desirable for estimation purposes to have as many bills as possible in the assessment, a surrogate "annualized"[1] consumption was constructed to represent the "pre" billing period.

For the post-installation "bills," the metered total from the quad-meter system was consolidated to make an annual bill for a period of at least a year. When the analysis data set was assembled in March of 2009, it was not possible to collect a year of billing data that would be clear of the installation date. Most of the DHPs were installed either in the fourth quarter of 2007 or the first quarter of 2008. A procedure similar to that used on the "pre" bills was used to construct a post-"bill" from the submetered data.

Table 2 is a raw comparison of before-and-after installation total kWh consumption for each of the 14 sites, uncorrected for any differences in weather from year to year. The "pre" figures are from monthly bills; the "post" figures are from hourly logged data. This data could be interpreted as the "raw" savings from the actual DHP installation. Missing values reflect the absence of whole-house data logging in two cases, and only six months of available pre-installation bills in a third. A minimum of a full year of "pre" bills was needed. The aggregate change, for the 11 sites with both figures available, shows roughly a decrease in electrical energy consumption of 2,800 kWh per household (12% decline). This change is very close to the median percent change, a decline of roughly 11%, and the median consumption decrease of 2,704 kWh.

It is important to note that these changes do not account for the changes in weather that occurred between the base years and the post-installation period. This analysis is essential to assess the space conditioning energy savings that could be attributed directly to the DHP installation.

[1] "Annualized" means that when more than a year of data was available, observations were weighted so that when the data were summed they represented only 365 days. For example, if logged data was available for a year and a half, days in the half of the calendar year which had two observations (from adjacent years) were each given a weight of .5, whereas the half-year of days with only one observation (from a single year) were given a weight of 1. All the assigned weights summed together, always equal 365.

Table 2: Total Whole House Annualized Consumption (kWh/yr)

	Consumption		
	Pre-bills	Post "Bills" 2008-2009	Raw Savings
Monmouth 1	26332	23062	3270
Monmouth 2	20836	19633	1203
Monmouth 3		11012	
Monmouth 4	22382	23723	-1341
Monmouth 5	16703	13733	2970
Monmouth 6	23451	19997	3454
Monmouth 7	20410	19193	1217
Monmouth 8	32631	28912	3719
Monmouth 9	20139	21171	-1032
Monmouth 10	37162	25102	12060
Monmouth 11	23829	21330	2499
MosesLake 1	12162		
MosesLake 2	29244		
Tacoma_1	15483	12779	2704
Average(of 11)	**23578**	**20785**	**2793**
Median(of 11)			**2704**

3.2. Weather Adjustment

The Table 2 summary does not take into account the differences in the weather between the pre-year and the post-year. Weather normalizing, or at least adjusting for differences in weather is desirable. Table 3 reports "pre" and "post" heating degree-day (HDD) comparisons for National Weather Service Cooperative Station Network sites that best matched the 14 house locations[2]. The 11 Monmouth, OR

---

[2]In comparison with the relatively sparse network of highly automated stations run by the Weather Service itself, the Cooperative Station Network is a dense network of stations, with hundreds of locations in the state of Oregon alone. However, not all the data are of high quality and not all stations had data records over the intervals of interest. We chose geographically close cooperative stations with recorded daily minimum and maximum temperatures available over the houses' "pre" billing period and "post" submetered data periods. The HDD to each base were calculated from daily minimum and maximum temperatures using the standard approximation formula: $Max(0, t_{base} - (t_{\min} + t_{\max})/2)$ where t_{base} is the chosen degree-day base temperature, and t_{min} and t_{max} are the day's recorded minimum and maximum temperatures, respectively.

sites had preliminary bills covering approximately the period between November 15, 2006, and November 15, 2007, and logged data covering slightly more than a year starting early February 2008. By contrast, the two Moses Lake sites had sufficiently different "pre" and "post" periods that we cumulated the Ephrata data twice, over the different intervals. There is a consistent pattern of increase in degree-days (increase in heating requirements) between the pre-installation period and the post-installation period. This increase occurs independent of the particular locality and independent of the calculation base from which the climate is evaluated. With Salem weather data, for example, the base 55º F HDD change increased 12.7%, but at base 65º F, the increase was only 7.1%. Although 65 º F is the conventional base for HDD calculation (based primarily on insulation levels of historic homes), HDD calculated to lower bases are, as we shall see, better predictors of temperature-dependent consumption (heating energy) for the homes in this study. The 11% pre-to-post median consumption declines seen in Table 2 occurred in the face of significantly colder weather in the "post" period, and savings estimates need to be adjusted to take this into account.

Table 3: Annualized Heating Degree-Day Comparisons

Weather station	Periods	Base 55º F	Base 60º F	Base 65º F
Salem (Monmouth)	11/15/2006-11/15/2007 Pre	2169	3370	4744.5
	02/08/2008-02/08/2009 Post	2444	3671	5082
	Pre-to-Post Percent Change	+12.7%	+8.9%	+7.1%
Ephrata (Moses Lake)	01/26/2003-09/25/2006 Pre	3329	4425	5689
	06/25/2007-09/23/2008 Post	4077	5224	6499
	Pre-to-Post Percent Change	+22.5%	+18.1%	+14.2%
Ephrata (Moses Lake)	10/06/2001-10/05/2004 Pre	3218	4334	5598
	06/05/2007-09/03/2008 Post	4074	5220	6511
	Pre-to-Post Percent Change	+26.6%	+20.4%	+16.3%
McMillin Reservoir (Tacoma)	12/13/2005-01/12/2007 Pre	2505	3790	5343
	02/15/2007-12/31/2008 Post	2716	4058	5632
	Pre-to-Post Percent Change	+8.4%	+7.1%	+5.4%

3.3. VBDD Energy Savings Analysis

To put the "pre" and "post" periods on an equal weather footing, we fit standard variable-base degree-day (VBDD) regressions to the "pre" metered bills for each of the 13 sites for which we have at least 12 months of "pre" bills. Where we had more than one year, the bills for all available months were included in the regression with the appropriate weather data. This step led to a more reliable fit to the temperature data and a better estimate of the heating energy before the DHP installation.

The VBDD regression methodology simultaneously estimates a house balance point (heating degree-day base), a slope coefficient of linear energy consumption response to heating degree-days, and a constant term which has an interpretation as unvarying monthly baseload (i.e., the sum of all non-space conditioning loads such as water heat and appliances). The "balance point" refers to the coldest temperature at which no space heating is required. The regression estimates this value and uses it as part of estimating the overall space-heating load.

Figure 2 displays a typical scatter plot that illustrates this analysis. This site (Monmouth 6) was analyzed comparing monthly kWh/day consumption (generated from electric bills) against degree-days per day (generated from the Salem, OR weather data for the pre-installation period) to balance point for that site

(59°F). The ascending straight line is the fitted regression line that captures the response of monthly kWh to heating degree-days. In fact, two separate lines are plotted, one with zero HDD months included, the other excluded. In this case, the exclusion has virtually no effect on the regression line, so the lines overlay one another and only one line is visible. Estimated coefficients and degree-day base (balance point) from these regressions provide a way to disaggregate billed consumption into heating (HDD-sensitive consumption) and "other." They also offer a way to predict heating consumption given the change in the weather data and a new set of temperature data. The R^2 for this site is typical of these homes and shows a good relationship between weather conditions and heating energy consumption. We applied the coefficients estimated using the "pre" period data to the weather data experienced in the "post" submetering period to estimate the hypothetical heating consumption that would have occurred in the "post" period without the DHP installation. Appendix A provides a more detailed explanation of this estimation procedure.

Figure 2: Point Scatter and VBDD Regression Lines for a Representative Site

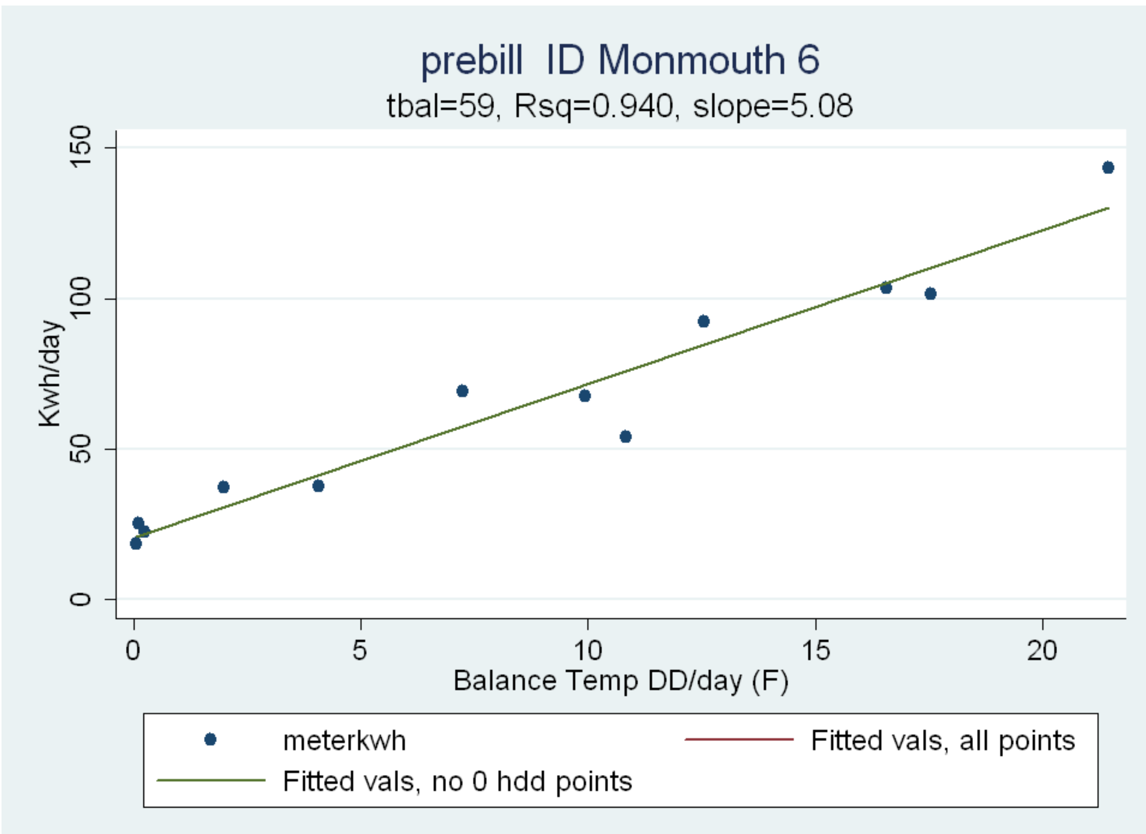

Table 4 summarizes the estimated space heat consumption from the billing VBDD analysis and the actual observed space conditioning consumption in the post-installation year. The "Pre-Bills Unadjusted" annualizes the estimates from the VBDD billing data regression estimates of the "Pre" annualized heating consumption. The "Pre-bills Weather-Adjusted" shows the results of adjusting the space heat estimates for the heating degree-days that occurred in the post-installation year. The "Post-Submetered" shows the accumulated space-conditioning consumption for the post-installation year and is directly comparable to the adjusted pre-bills calculation.

A relatively small but unknown proportion of the DHP consumption is cooling. Including it all is consistent with the conservative assumption that all cooling use of DHPs is new space-conditioning load; that is, none of it substitutes for preexisting air conditioner unit consumption. It would be useful to remember that space cooling remains a minor part of the annual energy use in most Pacific Northwest homes.

Cooling is relatively important in the Moses Lake area of eastern Washington. This distinction probably accounts for some of the reduced savings in that climate in comparison to the milder western Oregon and Washington climates. The presence of cooling reduces the apparent heating savings by increasing the submetered DHP usage. This effect is compounded since cooling could bias the VBDD by inflating the baseload estimate, thereby reducing the apparent heat load throughout the heating season.

Table 4: Annualized Heating Comparisons

Case	VBDD Estimated Annual Space Heat (kWh/yr)			Savings (kWh/yr)	
	Pre-Bills Unadjusted	Pre-Bills Weather-Adjusted	Post-Submetered (Resistance + DHP)	Unadjusted Heating Consumption	Weather-Adjusted Heating Consumption
Monmouth 1	8863	9917	6802	2061	3115
Monmouth 2	3265	4011	2356	909	1655
Monmouth 3			4557		
Monmouth 4	9071	12513	8893	178	3620
Monmouth 5	4243	5669	3952	291	1717
Monmouth 6	15140	17701	8081	7059	9620
Monmouth 7	5774	7018	5050	724	1968
Monmouth 8	14707	18561	10052	4655	8509
Monmouth 9	4574	5264	3797	777	1467
Monmouth 10	20429	24253	7246	13183	17007
Monmouth 11	5158	5391	2420	2738	2971
MosesLake 1	3341	4308	3540	-199	768
MosesLake 2	7665	10384	6902	763	3482
Tacoma 1	4485	4103	2259	2226	1844
Average(of 13)	**8209**	**10400**	**5537**	**2720**	**4442**
Average (of 11)	**8701**	**9930**	**5488**	**3164**	**4863**
Median (of 13)				**2061**	**2971**

The two "Savings" columns represent, respectively, the change in space-conditioning loads without weather adjustment, and the weather-adjusted change in space-conditioning loads. In this table, it is possible to make the relevant comparisons with all 13 sites, rather than 11 from Table 1, since the two sites which lack logged whole-house "post" consumption have logged DHP and resistance "post" consumption. We calculated averages with the 11 sites as well as the 13 to facilitate comparison with the Table 1 results. Table 4 summarizes aggregate results.

Table 5: Summary of Average Per-Site Annualized kWh Savings

	N=11	N=13
Unadjusted Consumption	2793	
Unadjusted Space Heat Only	3164	2720
Weather-Adjusted Space Heat	4863	4442

Conceptually, the weather-adjusted savings estimates presented in Table 4 and Table 5 are valid for the recorded post-installation period that happened to occur. To even the comparison playing field, we have used VBDD parameter estimation to shift pre-installation consumption patterns estimated from billing data to the post-installation period; we have not weather-adjusted post-installation submetered data. These savings estimates are not long-term-average weather savings estimates, nor are they synthetic “typical” TMY weather savings estimates. They are *weather-adjusted* savings, not *weather-normalized* savings.

4. Estimation Issues

4.1. Savings Outlier Sites

One savings outlier among the 13 sites, Monmouth 10, contributes disproportionately to average savings numbers. Weather-adjusted annualized savings for Monmouth 10 were estimated at over 17,000 kWh, almost 30% of the total savings estimated for this 13-site sample, and almost four times the site average. Normalizing by the 2800 sq-ft area of Monmouth 10, the calculated savings is over 6 kWh/sq ft/yr. Are these savings numbers credible? Appendix B examines this site in as much depth as our very limited characteristics data permit, in an attempt to assess the validity of this point. The conclusion is that the savings at this site are in fact large, but seasonally varying residual baseload, and/or possible contamination of submetered baseload by undetected resistance heating, are likely to be exaggerating the already significant savings number.

4.2. Cooling

The presence of summertime cooling load poses challenges for the VBDD methodology. One approach that has been employed in significant cooling climates is to complicate the VBDD model by adding cooling degree-days as an explanatory variable. This action adds two parameters to the three-parameter model—the cooling degree-day base, and the cooling degree-day response coefficient.

In practice, in Northwest cooling climates with transient, irregular, and small cooling loads this has never been very satisfactory. Window AC use is not very well described by cooling degree-days since they are not subject to a consistent thermostat setting. Because coefficients are typically ill behaved and unstable, we have not pursued this approach. Roughly half the sites had window AC units prior to the DHP installation, and we can assume that some cooling occurred in much of the sample, both before and after DHP installation. Prior bills, and post-installation logged data can provide clues that cooling load has occurred. In most of the sample, those clues suggest cooling load that is nonexistent or very small relative to annual heating load.

Nonetheless, it is possible in a handful of the 14 cases to detect a more significant pre-installation cooling load, and we examine those cases here. In the VBDD context, the obvious clue to cooling use is increased load occurring in zero HDD months, with billed usage minimums in shoulder season months rather than midsummer. Figure 3 depicts the fitted VBDD billing data regression for the Monmouth 2 site. Note the scatter of zero HDD points along the Y axis, mostly above the fitted lines, indicating evidence of anomalously high usage months without heating loads. Note also that two fitted regression lines—one which ignores those zero HDD points, and a second which includes them—have significantly different slopes. The "all points" line has a flatter slope, and a higher intercept, than the line that excludes zero HDD points. In the VBDD estimates of HDD response, we use the steeper green line, with the zero HDD points excluded, to remove the effect of the cooling load. The Figure 3 graph can be contrasted to the more typical VBDD regressions displayed in Figure 2 and Figure 6 (see Appendix B). In both those cases, there is no zero HDD point scatter along the Y axis, and the "all points" and "0 HDD excluded" regression lines are effectively identical.

Monmouth 2 is the only one of the 12 Westside climate zone 1 sites with an obvious cooling signature in billing data. The remaining 11 sites, even those with acknowledged window AC, do not display it. The two Grant county (Moses Lake) sites display something that looks like a cooling signature in pre-installation billing data. Figure 4 shows the fitted VBDD regression for the Moses Lake 2 site.

Figure 3: Monmouth 2 VBDD regressions

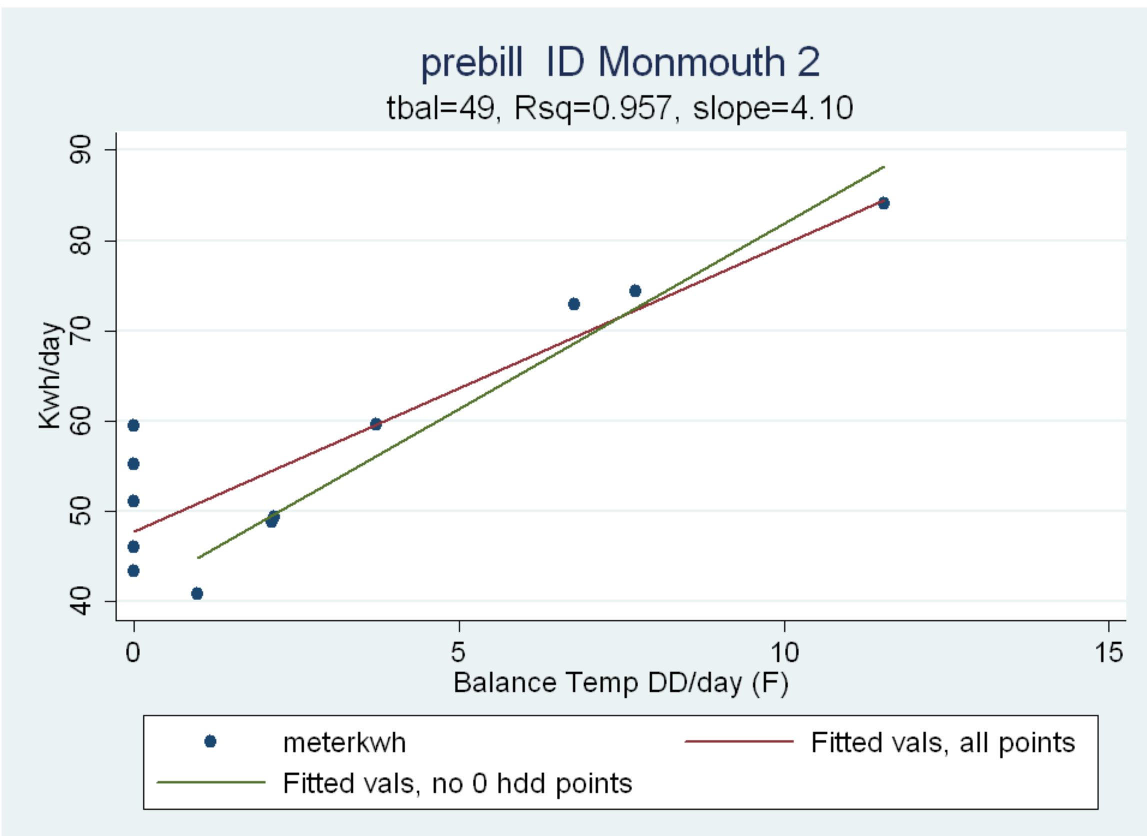

Figure 4: Moses Lake 2 VBDD regressions

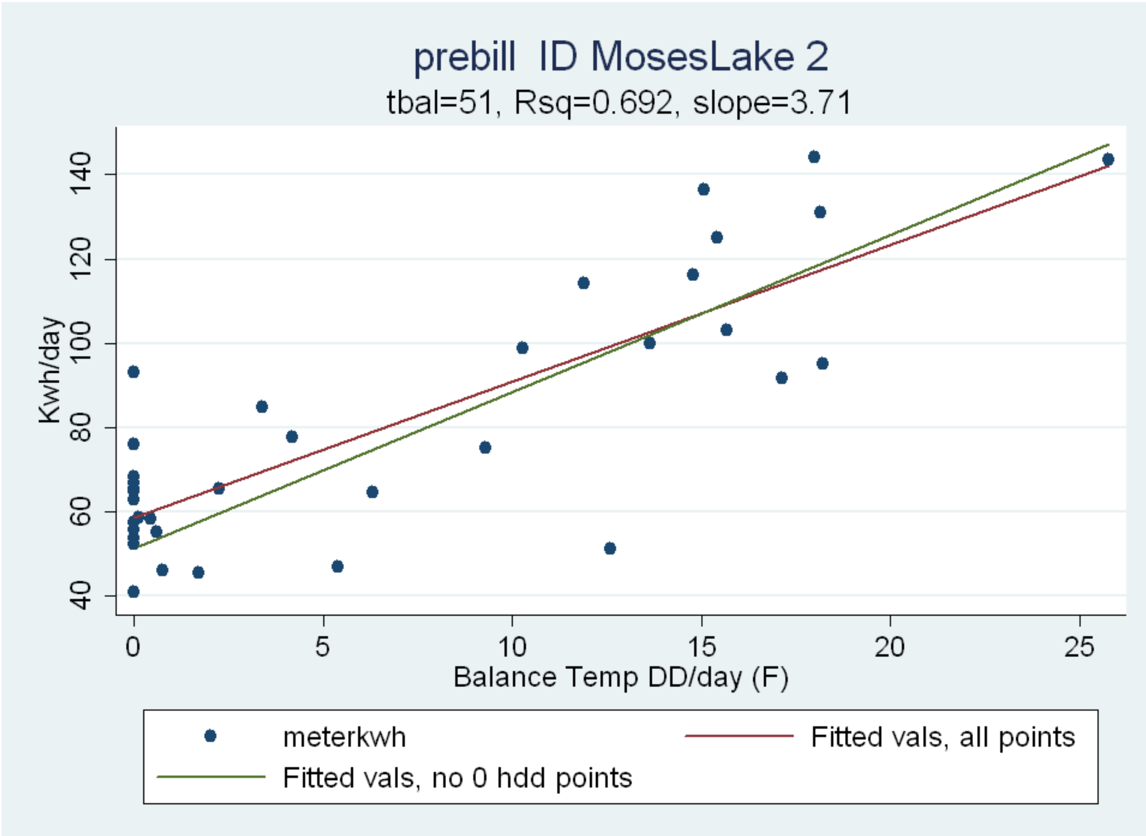

Cooling load also poses problems for interpreting the submetered logged data. After various experiments juxtaposing indoor and outdoor hourly temperature patterns with DHP usage, the conclusion is that there is no reliable method of assigning all hourly usage to heating or cooling without additional information such as vapor line temperature. Approximately 85% of positive usage hours can typically be assigned with reasonable confidence to heating or cooling, but the remaining 15% is ambiguous. Nonetheless, as with billing data, it is possible to detect approximate cooling loads in a few of our sites, based on seasonal usage. Appendix C displays graphs of monthly logged space-conditioning load juxtaposed with prior VBDD heating estimates, as well as graphs of all logged data streams, by month. In both graphs, it is possible to see what looks like DHP summer cooling in the sites that unambiguously had it before (Monmouth 2, Moses Lake 1, Moses Lake 2). The remaining 11 sites, which had no obvious cooling signature before, continue to show a pattern of minimal or no cooling.

There is no apparent evidence based on the current data that the installation of DHPs has unleashed significant new cooling load. Sites that had notable pre-installation cooling load evident in bills continue to have it in DHP usage; but the sites that did not have much cooling before, do not seem to use their DHPs much in cooling mode. More detail can, and should be added to this picture when we technical means are in place in the field to reliably distinguish DHP heating and cooling usage.

4.3. Cross-Sectional Relation of Savings to Characteristics Variables

In this section, we explore the variables that could be used to explain or generalize the savings results using the data available. This data is more limited as the occupant characteristics were not collected in the installation process. Nevertheless, the cross-sectional relationship between estimated weather-adjusted savings and available variables including floor area, estimated balance point, and prior total energy consumption was attempted. Table 6 displays the available data set. We are interested in any evident relationship between the first three numeric columns and the last one. With larger samples and more complete occupancy information, additional relationships could be explored. Even for this small sample, however, this simplified analysis is still illustrative.

Table 7 shows regression coefficients and significance levels for various simple and multiple regression specifications using the first three numeric columns of Table 6 as explanatory variables, and the weather-adjusted savings in the last column as the dependent variable. There are seven multiple regressions represented here. Thirteen observations (sites) are used in all the reported regressions. In a simple regression context, both prior annualized consumption (kWh) and estimated balance point (Tbal) have a significant positive relationship with weather-adjusted savings: the greater the prior total consumption, or the higher the balance point, the greater the magnitude of the savings. Floor area has a weak positive association with savings, but it falls short of standard statistical significance levels.

In a multiple regression context, prior consumption always retains strong significance. Balance point is significant when compared to floor area but not when it is included with prior consumption. Floor area is never significant in this data set even when used alone. Prior consumption and prior estimated balance point together explain close to two-thirds of the in-sample variability in weather-adjusted savings. Estimated regression constant term is usually significantly different from zero, and varies widely in magnitude from specification to specification depending on variables included. This pattern is not particularly interesting or significant. If all variables are mean-corrected before regression estimation, the response coefficients and R^2 would be unchanged, but the constant term approach zero.

Table 6: Cross-Sectional Summary

Site	Area (Sq Ft)	Pre-Consumption (from Table 2) (kWh/yr)	Estimated Balance Point (from VBDD) (°F)	Savings (from Table 4) (kWh/yr)
Monmouth 1	2600	26332	54	3115
Monmouth 2	1300	20836	49	1655
Monmouth 3	1800			
Monmouth 4	1650	22382	50	3620
Monmouth 5	1560	16703	56	1717
Monmouth 6	2025	23451	59	9620
Monmouth 7	2350	20410	53	1968
Monmouth 8	2100	32631	58	8509
Monmouth 9	3100	20139	53	1467
Monmouth 10	2800	37162	59	17007
Monmouth 11	1700	23829	57	2971
MosesLake 1	1100	12162	46	768
MosesLake 2	1250	29244	51	3482
Tacoma 1	1200	15483	53	1844
Average	**1903**	**23578**	**54**	**4442**

Table 7: Cross-Sectional Regression Coefficients for Seven Different Specifications (13 Sites)

Spec#	Constant	t-test	Area	t-test	kWh	t-test	Tbal	t-test	Adj R^2
1	-1359	-0.35	3.05	1.58					.11
2	-7805	-2.65**			0.53	4.30**			.60
3	-37642	-2.76**					783.79	3.09**	.42
4	-35617	-2.38**	0.80	0.43			717.66	2.35**	.37
5	-8243	-2.46**	0.50	0.33	0.51	3.50**			.56
6	-25255	-2.19*			0.40	2.81**	380.74	1.56	.64
7	-25637	-2.06*	-0.22	-0.14	0.41	2.62**	393.22	1.44	.60

Notes:

Spec#: regression specification number

Constant: regression constant

Area: floor area in square feet

kWh: prior annualized metered total kWh consumption

Tbal: house balance point estimated using VBDD regressions on pre-installation billing data

t-test: t-test statistic of the coefficient

Adj R^2: adjusted R-squared statistic for the regression specification

** : coefficient significantly different from 0 at 5% confidence level

*: coefficient significantly different from 0 at 10% confidence level

To interpret these results, it helps to keep in mind that the heat pumps were not sized to assume all the space-conditioning load in these houses, and in the case of the 11 Monmouth sites, identical units were installed in each of the houses, irrespective of house size. Effectively the DHP units themselves condition approximately the same amount of space in each house, with the remainder of the space (a greater proportion in larger houses) conditioned via other means. In this context, there is no reason to expect a larger house to achieve larger savings after heat pump installation.

The positive and significant relationship between total prior consumption and savings should be interpreted as a utilization effect. For a heat pump to save energy relative to the prior heating system, it has to be used. These higher energy savings are more likely to occur if the house and its occupants were already heavy energy consumers, and if (due to whatever combination of house UA, high thermostat set point, or low internal gains) the prior balance point is also high.

One should not read too much into these results. We have only 13 usable sites, and in some individual cases, the savings that are the dependent variable seem suspect. As we have stated, richer characteristics data would better enable us to disentangle real savings from extraneous changes, or from estimation, data, or modeling problems. As a robustness check on these results in the face of these unknowns, we re-ran the regressions from Table 7 with only 12 observations, excluding the biggest savings outlier site, Monmouth 10. For brevity, we have not reported the coefficients here. Although, dropping this site does not change the story much. In the 12-observation regressions, floor area has no relation with savings, even in a simple regression. Balance point and prior total consumption retain strong predictive power in a simple-regression context.

Figure 5 illustrates the point scatter and fitted line for the Table 7 regression specification #2, which uses only prior annualized total kWh as an explanatory variable. Prior total kWh is the explanatory variable with the strongest individual effect on calculated savings. The point from the savings outlier site Monmouth 10 is labeled “Mon 10.” As can be seen, the strength of the regression is heavily influenced by the outliers. Removing them would help the R^2.

Figure 5: Savings and Prior total Kwh

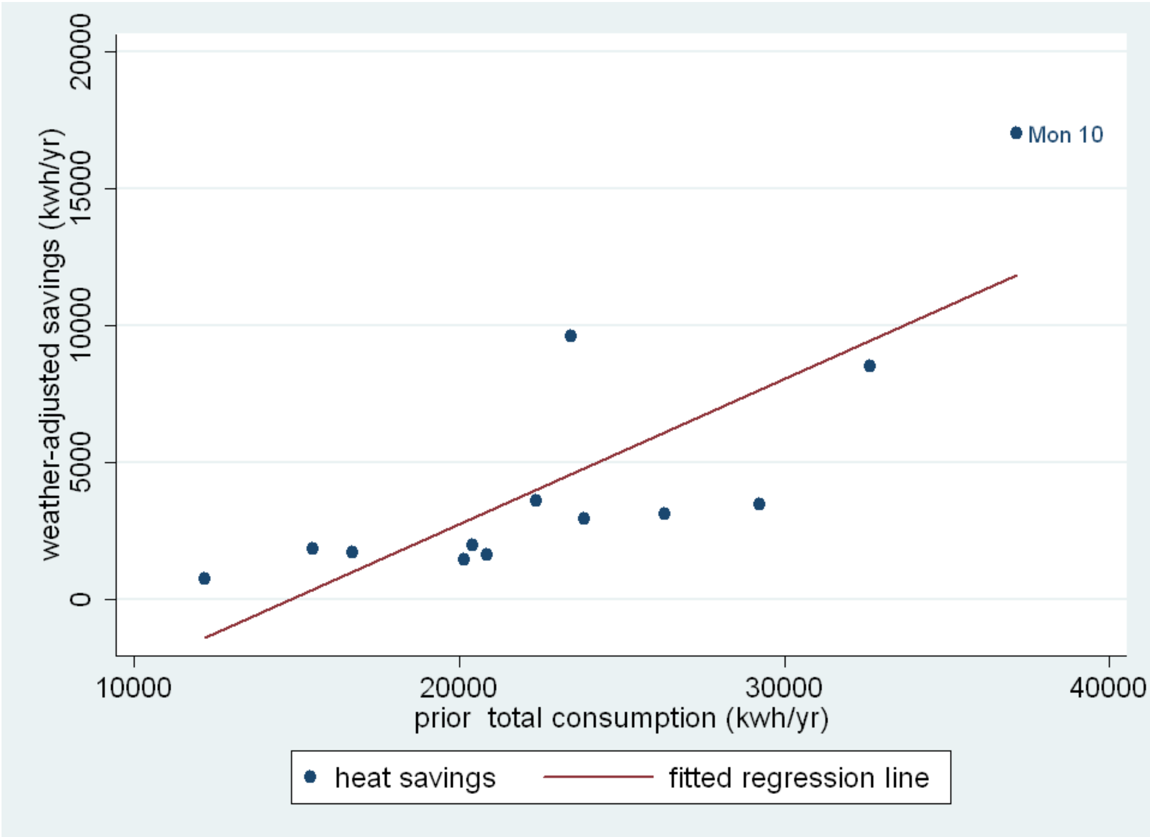

5. Conclusions

Leaving aside the nature of this study as a series of case studies without a strong relationship to any population, this review has provided many valuable insights and hypotheses that can influence future work:

- The most significant insight was identification of the metering package required to provide an effective data set for understanding the performance of the DHP systems. Because of the accessibility of thermostat set points in these installations it is very difficult to determine when the heat pump is in cooling mode. This problem was thought to be easily resolved by the use of outdoor temperature. It has proven ambiguous at a level that makes any estimate of cooling offsets problematic. The result has been the introduction of vapor line temperature to act as a surrogate for the cooling signal.
- The fact that these installations do not have metering before the installation of the DHP system means that the savings and changes in electric consumption depend on a billing analysis in the year prior to the installation. This analysis has been assisted by the use of more than one year of pre-consumption data. As a result, these homes have consistently high R^2 and reasonable estimates of space heating prior to installation. Due to the lack of sufficient time between the end of the metering period and the beginning of the analysis, the data set as collected did not include the utility consumption after the installation. Post-installation data would have been helpful in verifying the assumptions drawn from the meter results.
- There were fairly minimal characteristics data available. As with the billing data more characteristics data may have explained some of the details of the metered and billing data used here. It is our recommendation that enough data be collected to discern occupant patterns and habits as well as sufficient data on the home itself to estimate the underlying heating requirements. This would allow more validation of the base case billing analysis results and more direct simulation of the impacts of the DHP installation.

While there are important caveats that should be placed on this analysis there are several observations that can be thought of as preliminary indication of the impact of DHP installation using the protocol designed for this project:

- Initial saving estimates suggest that the savings associated with this technology is approximately 4,500 kWh per year. This estimate compares favorably with the estimates used by the Regional Technical Forum in assessing this technology as an energy savings measure.
- The use of cooling in this group does not seem to be increased by the introduction of the DHP technology. In general, the existing zone cooling system seems to be offset by the DHP or in the absence of a cooling system only minimal cooling energy increases from the DHP were observed.
- The use of the DHP as a displacement heating system has the impact of generating savings independently of initial house size. Since this strategy provides a uniform capacity in a single zone the size of the remaining zones in a particular house are less important. Presumably, the overall efficiency of the house envelope will have an impact on the savings estimates but that hypothesis will require more characteristics data and a more representative sample.

Many of these observations will be reviewed in the future research on this technology. At this writing the DHP systems reviewed here provide a promising indication of a viable future savings measure for existing electric zonal heated homes.

6. References

Fels, M., J. Rachlin, and R. Socolow. 1986. Seasonality of Non-Heating Consumption and Its Effect on PRISM Results. Energy and Buildings, Volume 9 (1986), pp. 139-148.

Fels, M. 1986. PRISM: An Introduction. Energy and Buildings, Volume 9 (1986), pp. 5-18.

Roos, C. and D. Baylon, 1993. Non-Space Heating Electric Consumption in Manufactured Homes, RCDP Cycle II, Ecotope, Inc. Seattle, WA.

Quaid, M., R. Kunkle, and B. Lagerberg, 1991, RCDP Cycle II Appliance Analysis, Washington State Energy Office, Olympia, WA.

Appendix A: Variable-Base Degree-Day Regressions

The study regressed billing period consumption on billing period degree-days using a slight modification of the standard variable-base degree-day method pioneered by Fels (1986). Under the Fels PRISM method, also known as variable-base degree-day (VBDD) regression, the heating degree-day base and the regression response coefficient of energy consumption to degree-days are jointly estimated by finding the heating degree-day base which maximizes "goodness of fit" as measured by R^2, the regression coefficient of determination. Using R^2 as a criterion effectively maximizes the proportion of total variation in consumption explained by a linear response to heating degree-days. In a single zone structure (like a manufactured house) heated with an electric resistance furnace and a seasonally unvarying baseload, the linear coefficient has the interpretation of house UA, and the regression intercept has the interpretation as a seasonally constant average baseload not dependent on space heating demand. The degree-day base estimated by this procedure has an interpretation as the house balance point. Balance point is not thermostat set point, but rather is the lowest outside temperature at which the set point temperature can be maintained without space heating—where house internal and solar gains precisely match heat loss. Except in the special and implausible case where house internal and solar gains are zero, balance point is lower than thermostat set point. Although 65º F is a plausible thermostat set point, it is not a reasonable balance point for the vast majority of houses. Varying solar gains and thermostat set point changes have the effect of changing the balance point, so that the actual heating input data (the bills) in fact reflect some random mix of effects of heating degree-days to different bases.

The "Ecotope modification" to the Fels PRISM procedure involves excluding data points from a regression estimation where the billing interval's heating degree-days to that base are zero. Empirically, this serves to insulate the estimated HDD slope coefficient and constant from the influence of summertime cooling loads, which certainly exist for some of our sites.

Given a variable-base degree-day (VBDD) fitted regression coefficient and estimated balance point, a straightforward estimate of heating load for a given month is the product of the regression coefficient with HDD to that balance point base for that month. An accompanying estimate of annual non-heating related base load is simply the fitted regression constant times 12 months. A problem with this simplest of approaches is that it is well established from submetered data that non-space-heat load components do have seasonal variation, notably electric light (with length of day) and hot water heat (with seasonally varying intake water temperature), and without adjustment these seasonally varying base load components are imputed to heating load. An adjustment method first proposed by Fels et al (1986) is to fit a cosine function using the regression constant. Following the Fels approach, we adjust our heating estimate using a trigonometric function of the estimated regression "base load" constant α as follows:

$$\text{Heat for month } m = Max(\beta \cdot HDD - \alpha \cdot (.1 + .1 \cdot \cos(2\pi m / 12)), 0)$$

Where β is the estimated regression slope coefficient, *HDD* is calculated heating degree-days for month *m* to the chosen base, and α is the estimated regression constant. In effect, some of the seasonally varying load is taken away from the heating estimate $\beta \cdot HDD$ and given to the base load estimate α.

Given estimated coefficients, the above formula can be used to predict heat consumption given a new set of HDD data—not the HDD data which were used in the actual coefficient estimation. This is how we derive our estimates of the heating consumption that would have occurred in the "post" period had the old heating system not been replaced by a DHP. The parameters estimated in the "pre" period are applied to the "post" period's HDD in the above formula. Although external temperature is one of our post-installation submetered data streams, and could optionally be used as a basis for post-installation period HDD calculation, we chose to continue with the same cooperative weather station temperature data stream that was used to estimate the “pre” billing data regressions.

Appendix B: Examination of Savings for Outlier Site Monmouth 10

Calculated weather-adjusted savings for Monmouth 10 were over 17,000 kWh/year, almost 30% of total estimated savings for 13 sites. The floor area of Monmouth 10 is about 2800 square feet, about 47% larger than the study group's average (but not the largest home). Unadjusted, annualized pre-installation total consumption at Monmouth 10 (from bills) was over 37,000 kWh, with an unadjusted drop in house total consumption of over 12,000 kWh, or 4.3 kWh/sq ft/yr. Remembering that there were significantly more heating degree-days in the post-installation period—roughly 10% more at a degree-day base of 59 degrees (the "best-fit" base for this house)—it is not a surprise that weather adjustments would make this raw savings figure appear larger still, but it was not expected that the adjustment be almost 50%, roughly from 12,000 to 17,000 kWh/yr. It is worth scrutinizing the adjustment steps in this case.

Figure 6 presents the data used in the VBDD analysis for Monmouth 10, and also displays the fitted regression line that was subsequently used to separate the total billed consumption into an HDD-dependent portion, assumed to be space heat, and a non-HDD-associated "baseload." It should be apparent that (like most of the fitted VBDD regressions for the 13 sites that received them) the R^2 is quite high and the point scatter around the regression line is narrow and homogenous, with no apparent nonlinearities or heteroskedasticity issues. Hence, we should have reasonable confidence in our disaggregation of the monthly-billed consumption figures into baseload and HDD-dependent portions.

Figure 6: Monmouth 10 VBDD Regression

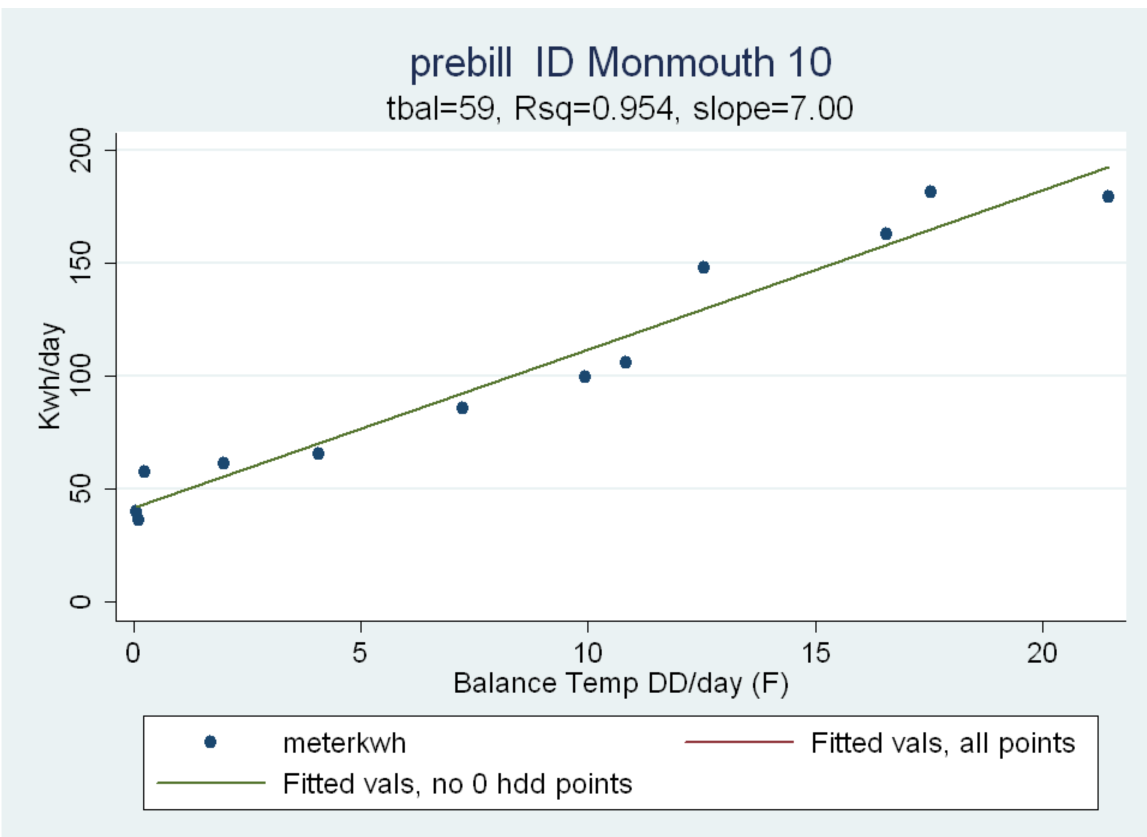

When we disaggregate the pre-installation billing data in this way, we find that about 55% of the consumption—slightly over 20,000kWh—is HDD-dependent. Comparing this figure with the annualized post-installation submetered space-conditioning consumption (DHP and resistance) we find a drop of roughly 13,000 kWh, slightly more than the raw decrease in consumption totals. The gap widens significantly, when the weather correction is applied to the parameters estimated in the pre-installation period. The weather correction indicates a pre-installation annualized heating consumption, adjusted for the greater HDD value in the "post" period, of 24,253 kWh/year. Normalizing this by square feet yields annual space heating of 8.66 kWh/sq ft. The post-installation annualized heating figure, from submetered DHP and resistance data, is 7246 kWh/year, or 2.59 kWh/sq ft. Both of these figures are individually plausible, but it is unlikely that installation of an (undersized) heat pump alone will effect such a change. It would imply an effective COP of about 3.3 if the DHP performed 100% of the post-installation space-heating task; and we know, in fact, that a significant portion of the post-DHP installation heating load in this house was borne by electric resistance.

One partial explanation for these improbable savings numbers is immediately suggested by Figure 7, a graph of submetered consumption for Monmouth 10 in the post-installation period. Submetered baseload—that is, the residual consumption after DHP usage, resistance space heating, and water heating consumption are subtracted—is strikingly seasonal. In fact it is so much so that it suggests that there is hidden or un-accounted for space heat consumption in this residual baseload. But even if this is not hidden space heat—even if there is some other strong persistent seasonal pattern in non-heating electricity use, both before and after HDD installation, it's still going to cause problems, because it is positively correlated with HDD and so will bias upward the VBDD regression coefficients estimated for the pre-installation period. Appendix C contains similar graphs of submetered data for all sites, displayed jointly, and it can be immediately appreciated from these Appendix C graphs that Monmouth 10 has an unusually strong winter-high residual baseload peak.

Figure 7: Submetered Monthly Consumption for Monmouth 10

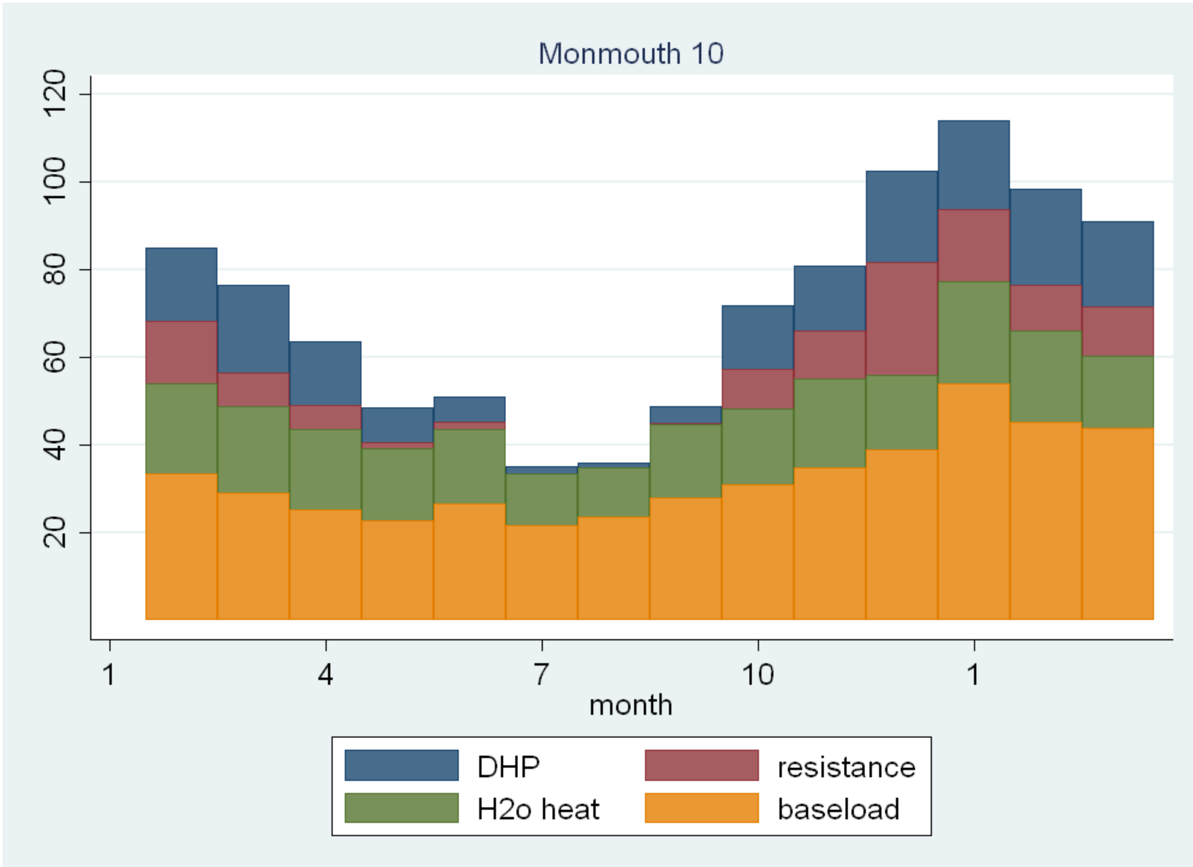

So either we have space heating hidden in our post-installation submetered baseload, or we have HDD-correlated seasonally varying baseload contaminating and biasing upward our VBDD estimates of heating energy response to HDD. Either way, our comparisons of space heating components, before and after the DHP installation, are suspect.

These potential problems of breaking total consumption into its space-conditioning and non-space-conditioning components, either before DHP installation (with VBDD regressions) or after (with submetering) do not change the large 12,000+ kWh drop in raw total annualized consumption between "pre" and "post" periods. In addition, the likelihood is that some unknown portion of that 12,000 to 17,000 kWh jump in estimated savings induced by our disaggregation and weather adjustment procedures is real. We are left with the conclusion that "something else" is probably going on here. It is possible that other behavioral or house characteristics changed coincidentally with DHP installation or that the installation itself was different from the other cases. Unfortunately, the lack of secondary data in this data set leaves these questions unanswered.

Seasonally varying baseloads represent one estimation problem that detailed site characteristics data would not solve. We conjectured above that our post-installation submetered baseload figures may be contaminated by hidden space heating. If that is the case, collection of detailed characteristics data would probably have brought it to light and avoided the problem; but if submetered baseload unrelated to space conditioning or water heat is seasonally highly variable that poses problems for VBDD regressions applied to billing data. This problem has long been recognized, and in the VBDD regressions we employ a trigonometric adjustment that in its conception dates all the way back to Fels (1986). The problem with this procedure, quite clear in the case of Monmouth 10, is that this is a one-size-fits-all adjustment and the magnitude and nature of baseload seasonality varies from site to site. A definitive fix to this issue—submetering for a pre-installation year, thus avoiding VBDD billing data regressions altogether—is not "in the cards," but we believe that site-specific secondary audit information may reduce this problem to a manageable set of adjustments using engineering methods.

Appendix C: Individual Site Graphs

The graphs included in this appendix represent individual site graphs of monthly, submetered data and before-and-after heating consumption comparisons.

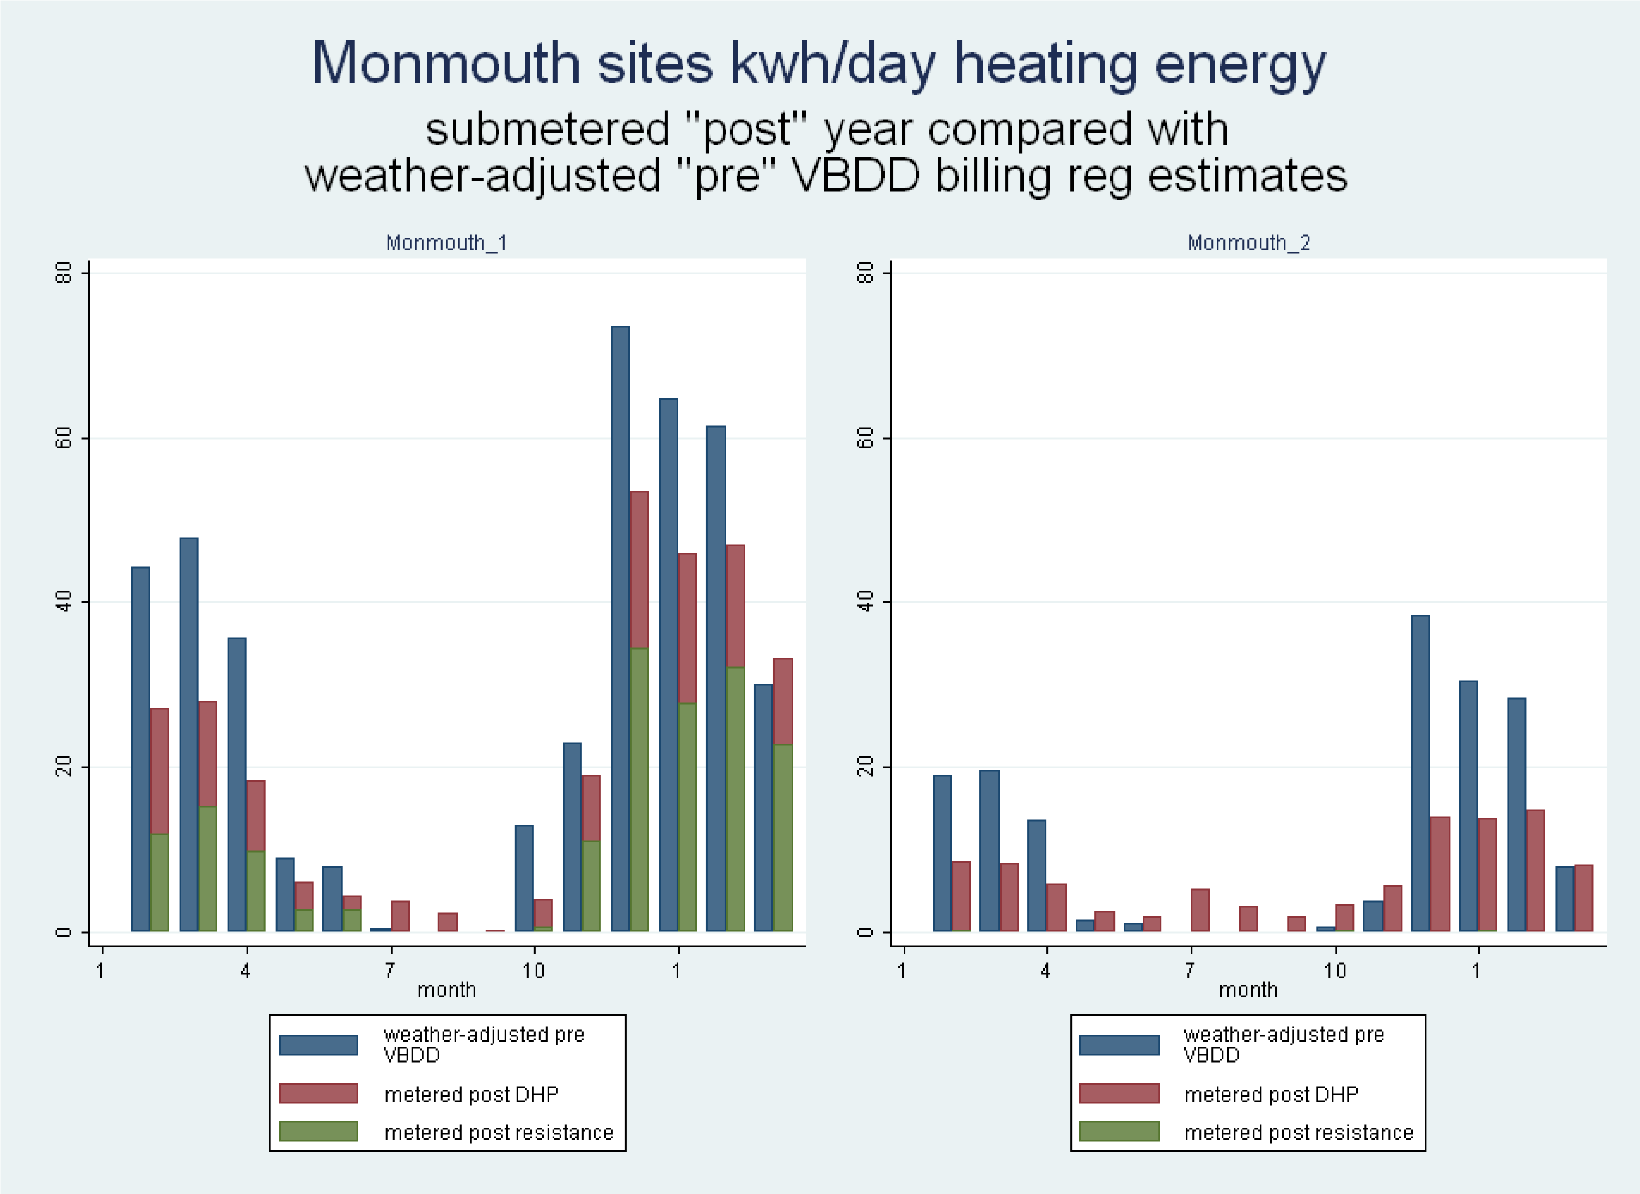
Monmouth sites kwh/day heating energy
submetered "post" year compared with
weather-adjusted "pre" VBDD billing reg estimates
Monmouth_1
Monmouth_2
0
20
40
60
80
1
4
7
10
1
month
weather-adjusted pre VBDD
metered post DHP
metered post resistance

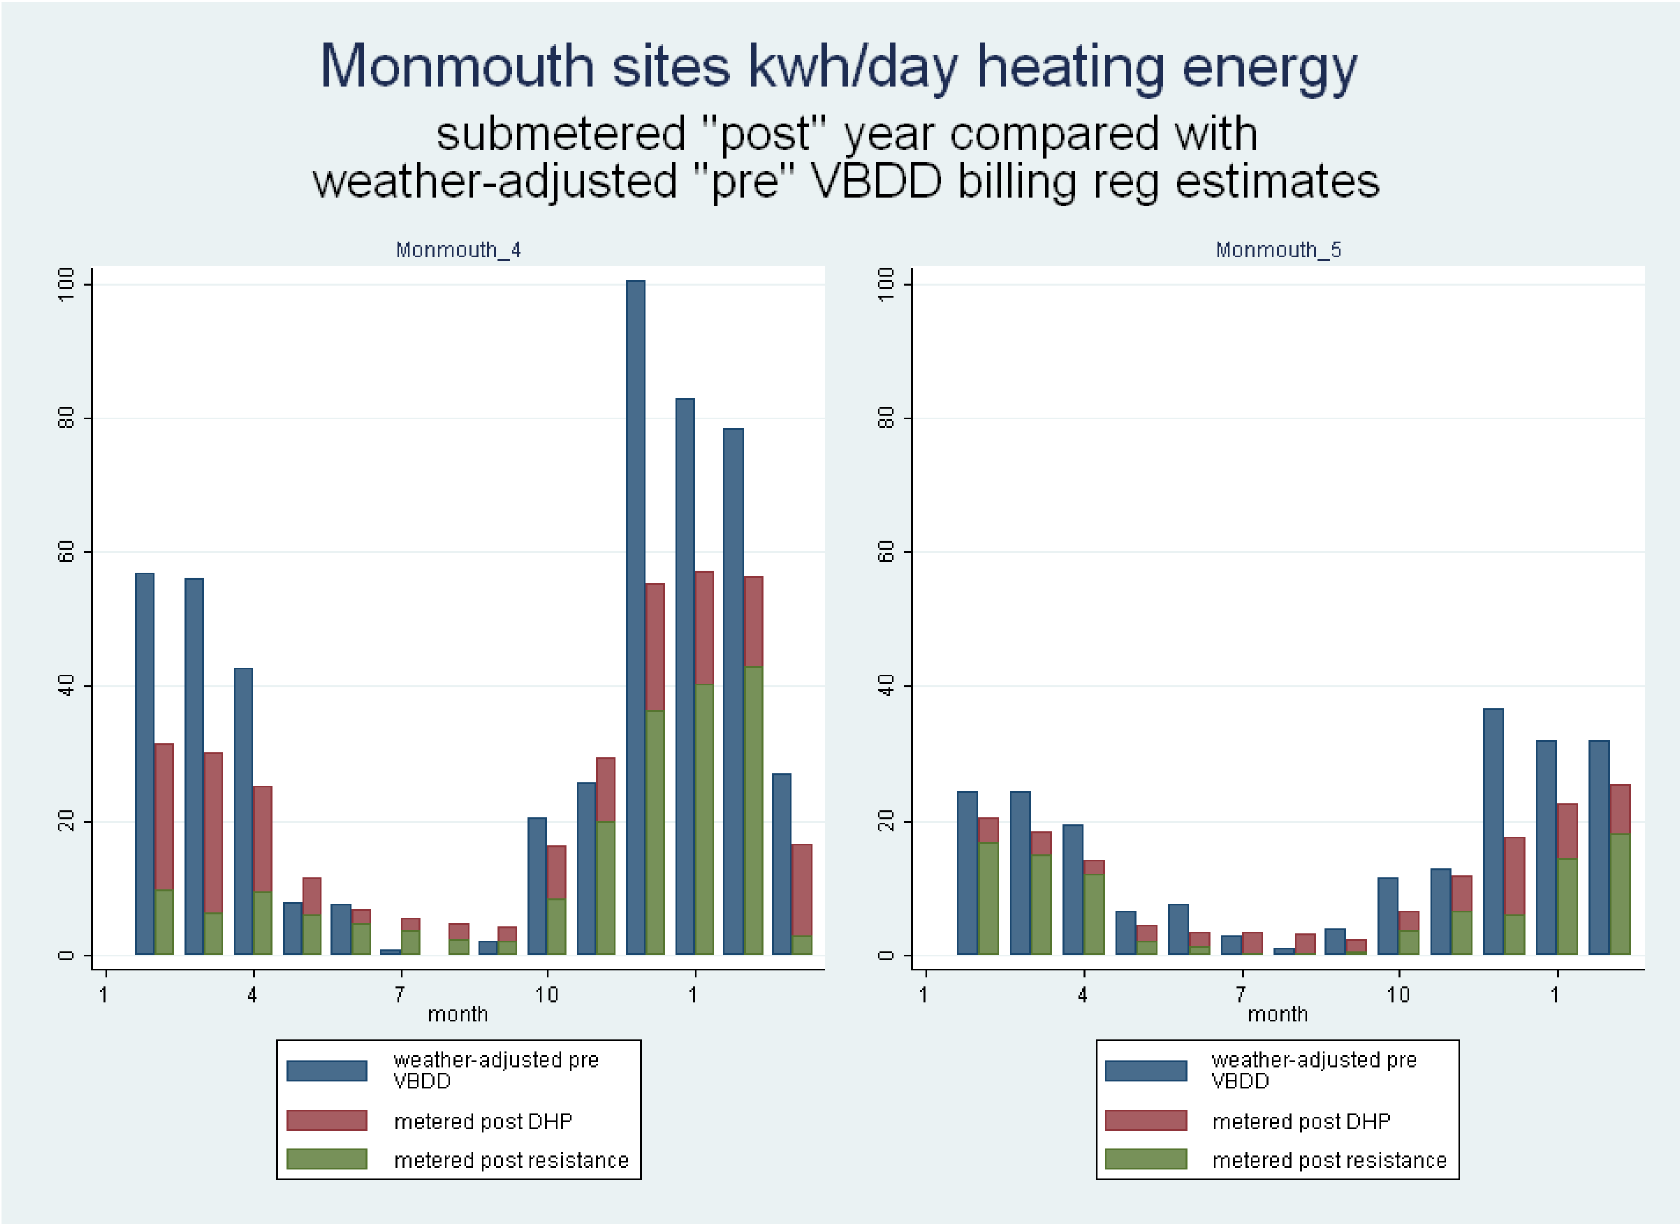

Monmouth sites kwh/day heating energy
submetered "post" year compared with
weather-adjusted "pre" VBDD billing reg estimates
Monmouth_4
Monmouth_5
0
20
40
60
80
100
1
4
7
10
1
month
weather-adjusted pre VBDD
metered post DHP
metered post resistance

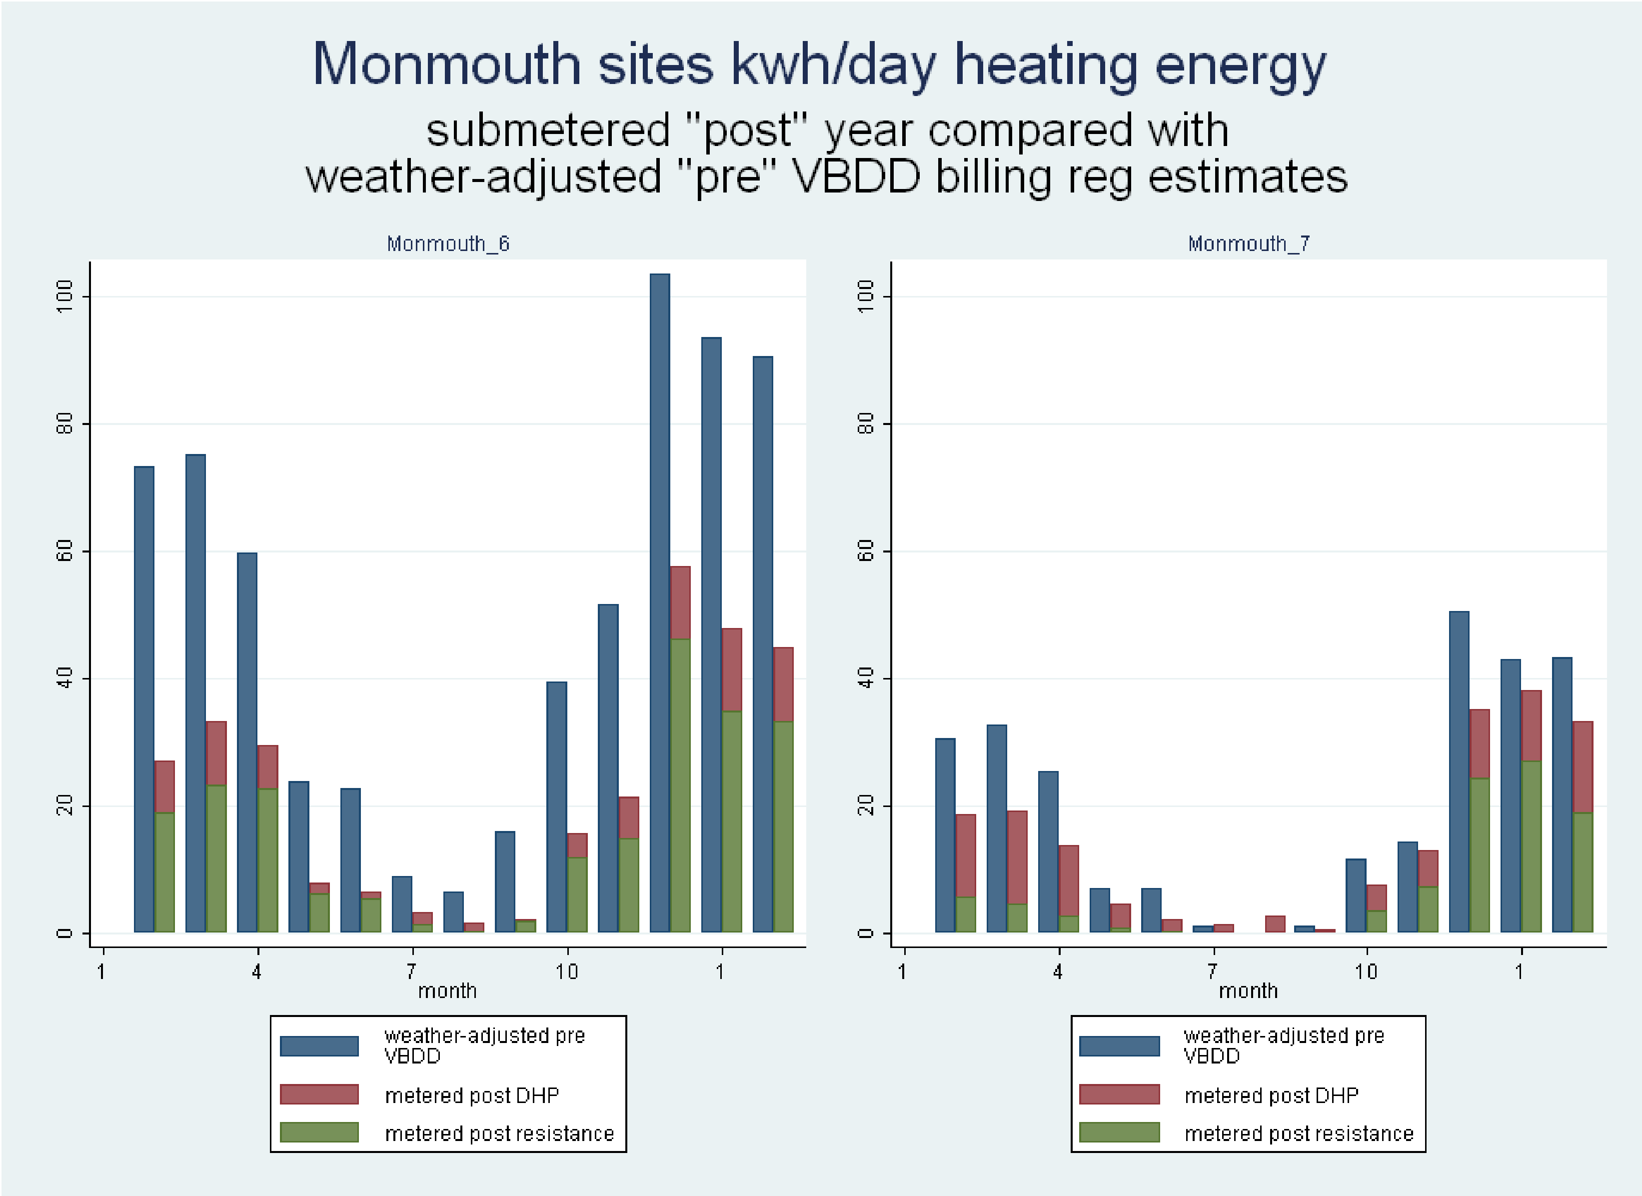
Monmouth sites kwh/day heating energy
submetered "post" year compared with
weather-adjusted "pre" VBDD billing reg estimates
Monmouth_6
Monmouth_7
0
20
40
60
80
100
1
4
7
10
1
month
weather-adjusted pre VBDD
metered post DHP
metered post resistance

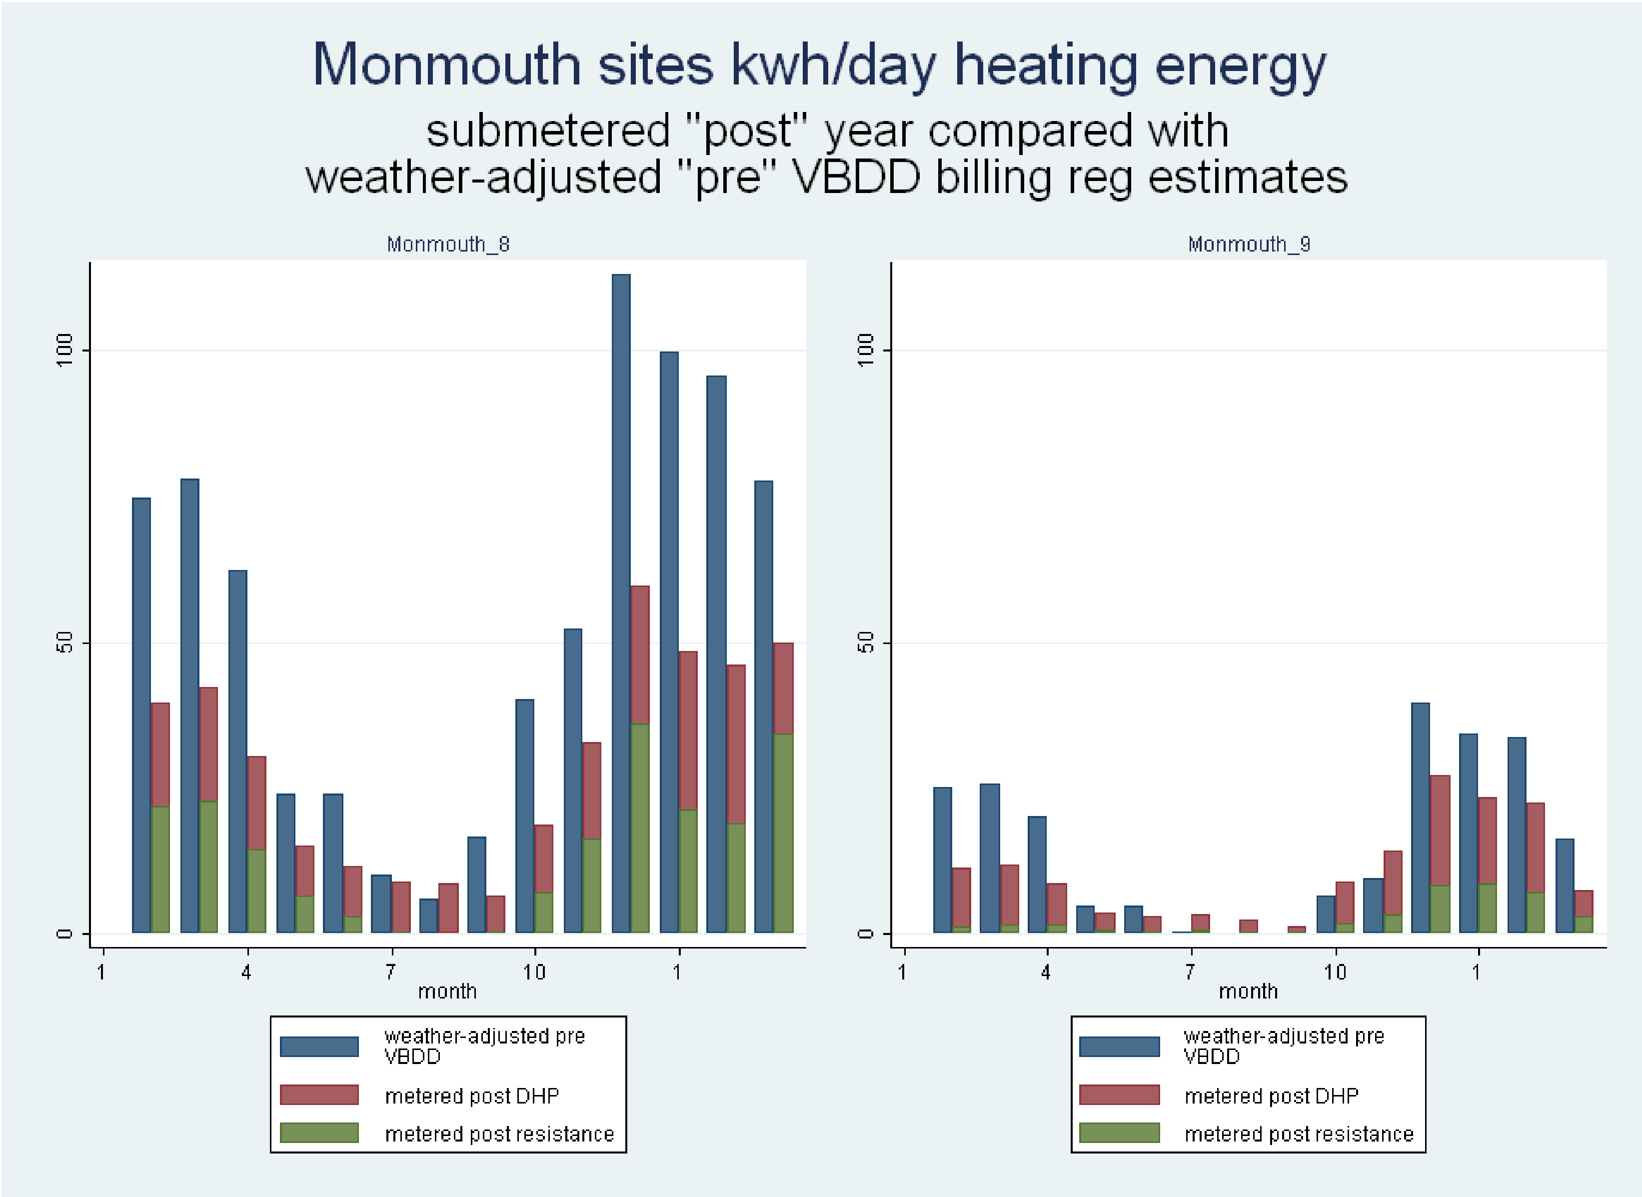
Monmouth sites kwh/day heating energy
submetered "post" year compared with
weather-adjusted "pre" VBDD billing reg estimates
Monmouth_8
Monmouth_9
0
50
100
1
4
7
10
1
month
weather-adjusted pre VBDD
metered post DHP
metered post resistance

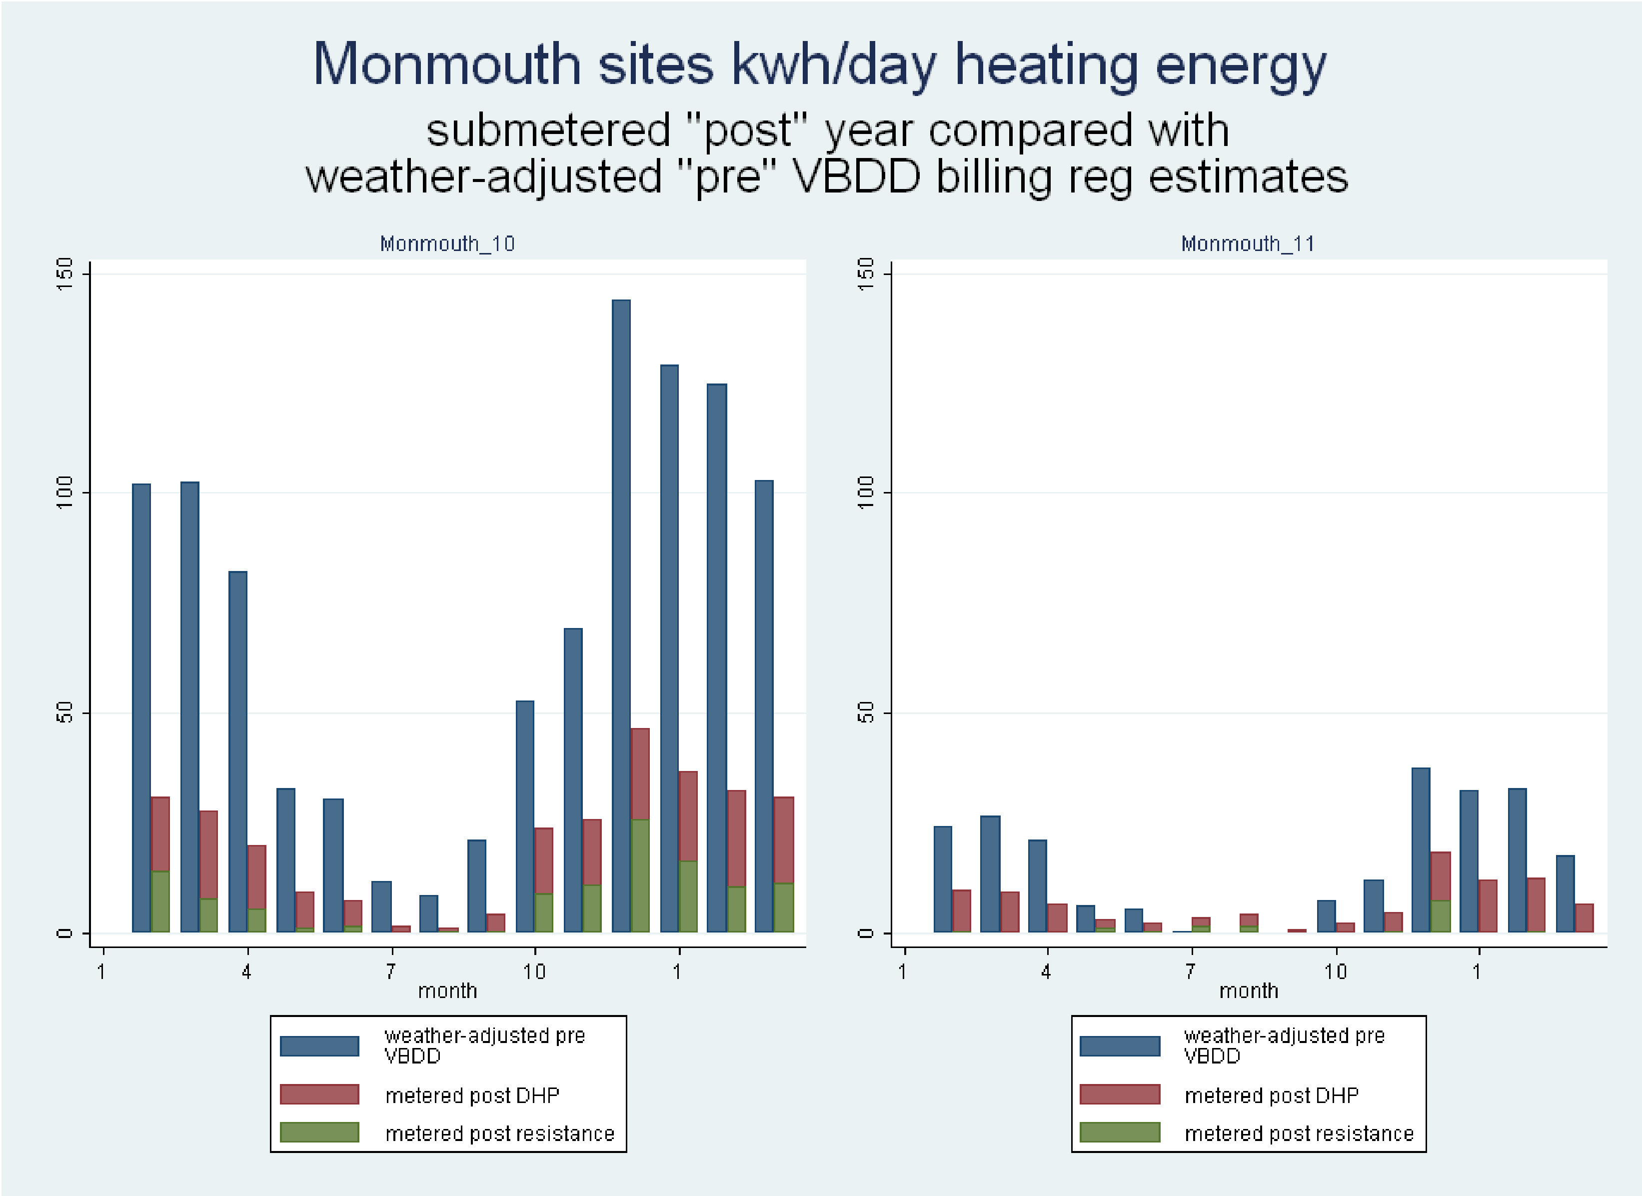
Monmouth sites kwh/day heating energy
submetered "post" year compared with
weather-adjusted "pre" VBDD billing reg estimates
Monmouth_10
Monmouth_11
0
50
100
150
1
4
7
10
1
month
weather-adjusted pre VBDD
metered post DHP
metered post resistance

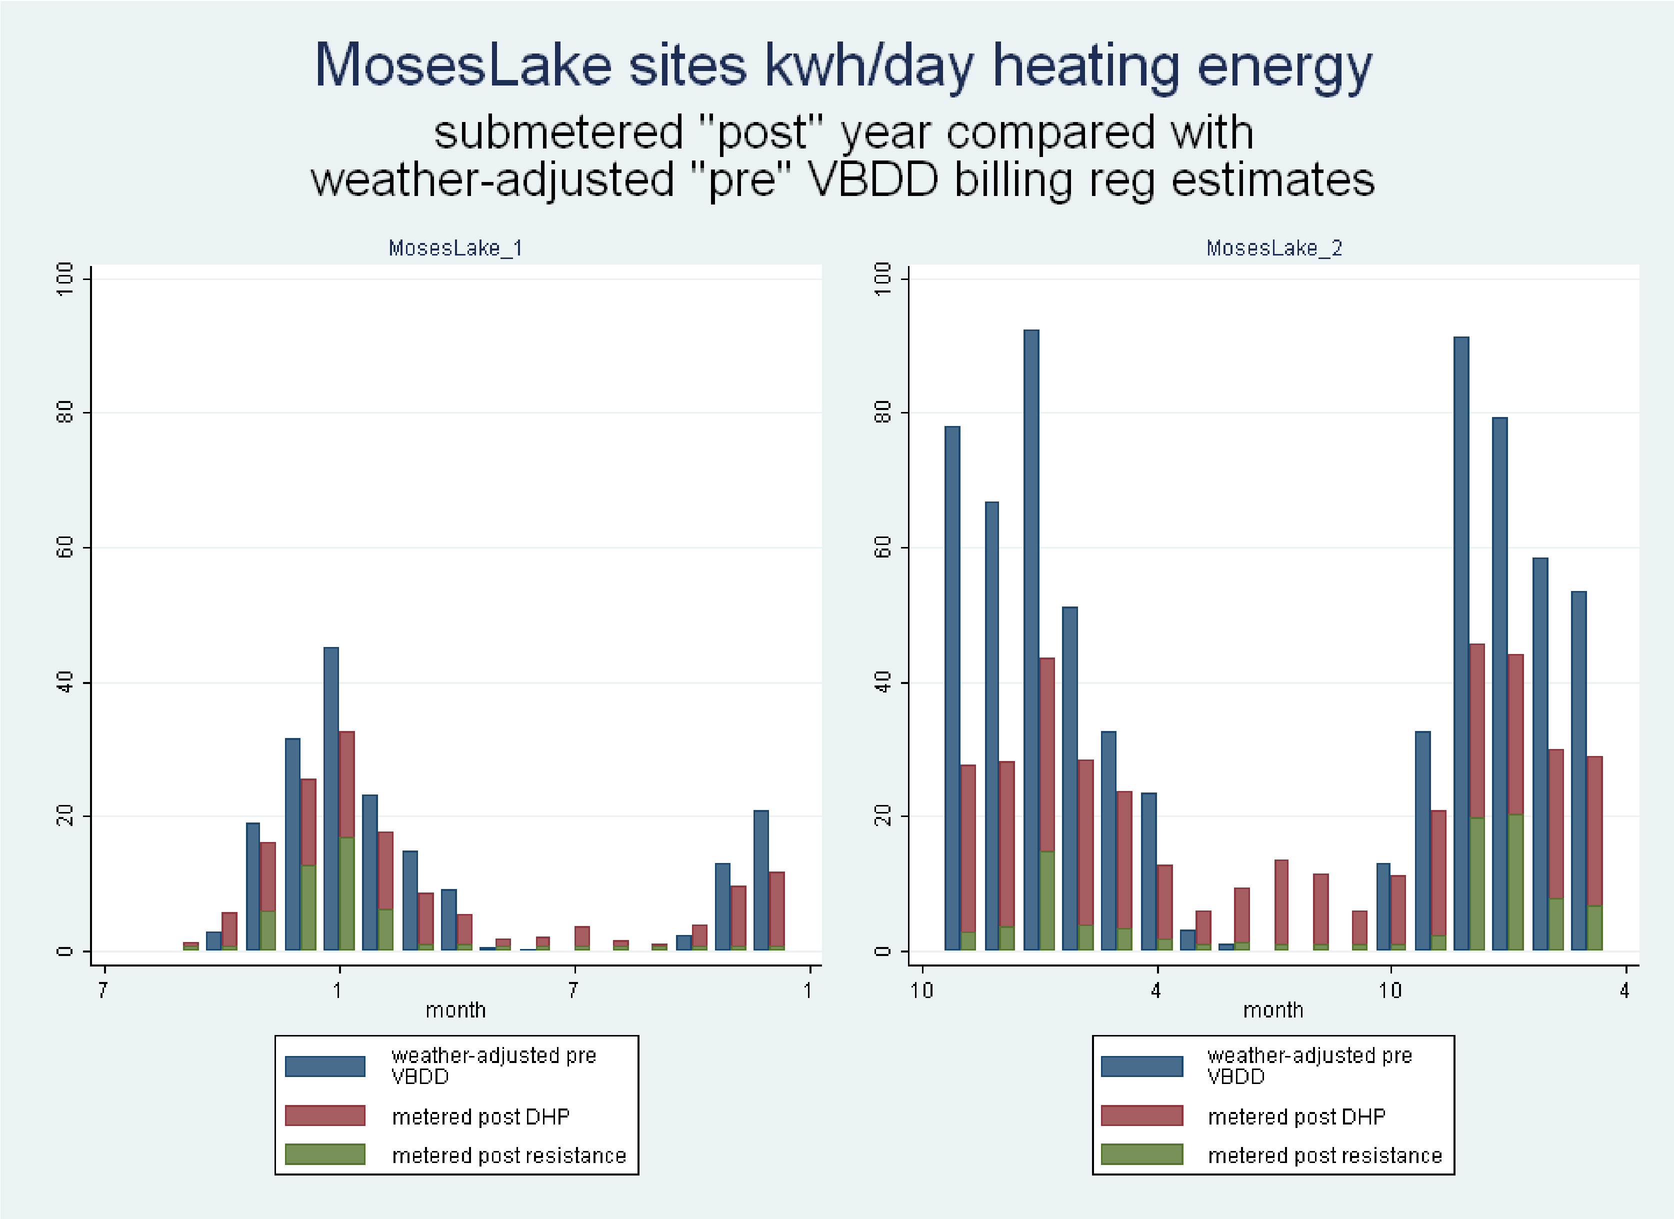
MosesLake sites kwh/day heating energy
submetered "post" year compared with
weather-adjusted "pre" VBDD billing reg estimates
MosesLake_1
MosesLake_2
0
20
40
60
80
100
7
1
7
1
10
4
10
4
month
month
weather-adjusted pre VBDD
metered post DHP
metered post resistance

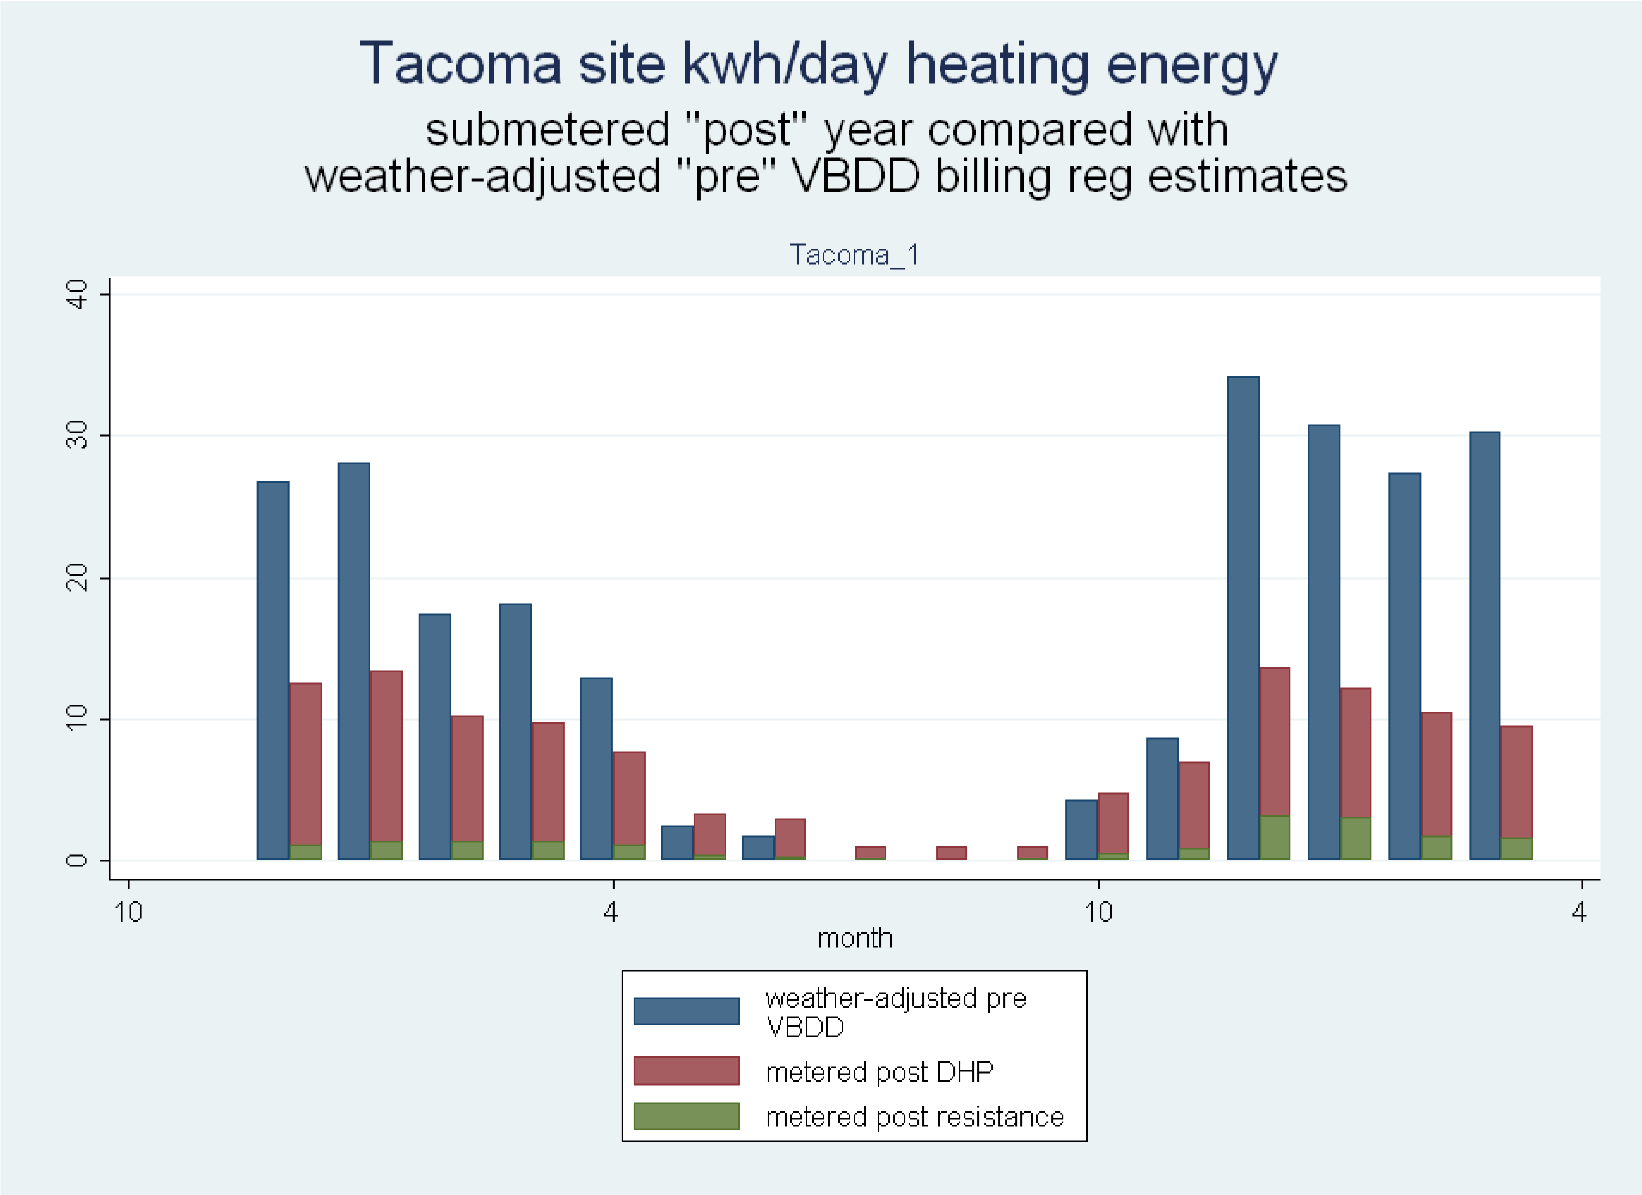
Tacoma site kwh/day heating energy
submetered "post" year compared with
weather-adjusted "pre" VBDD billing reg estimates
Tacoma_1
0
10
20
30
40
10
4
10
4
month
weather-adjusted pre VBDD
metered post DHP
metered post resistance

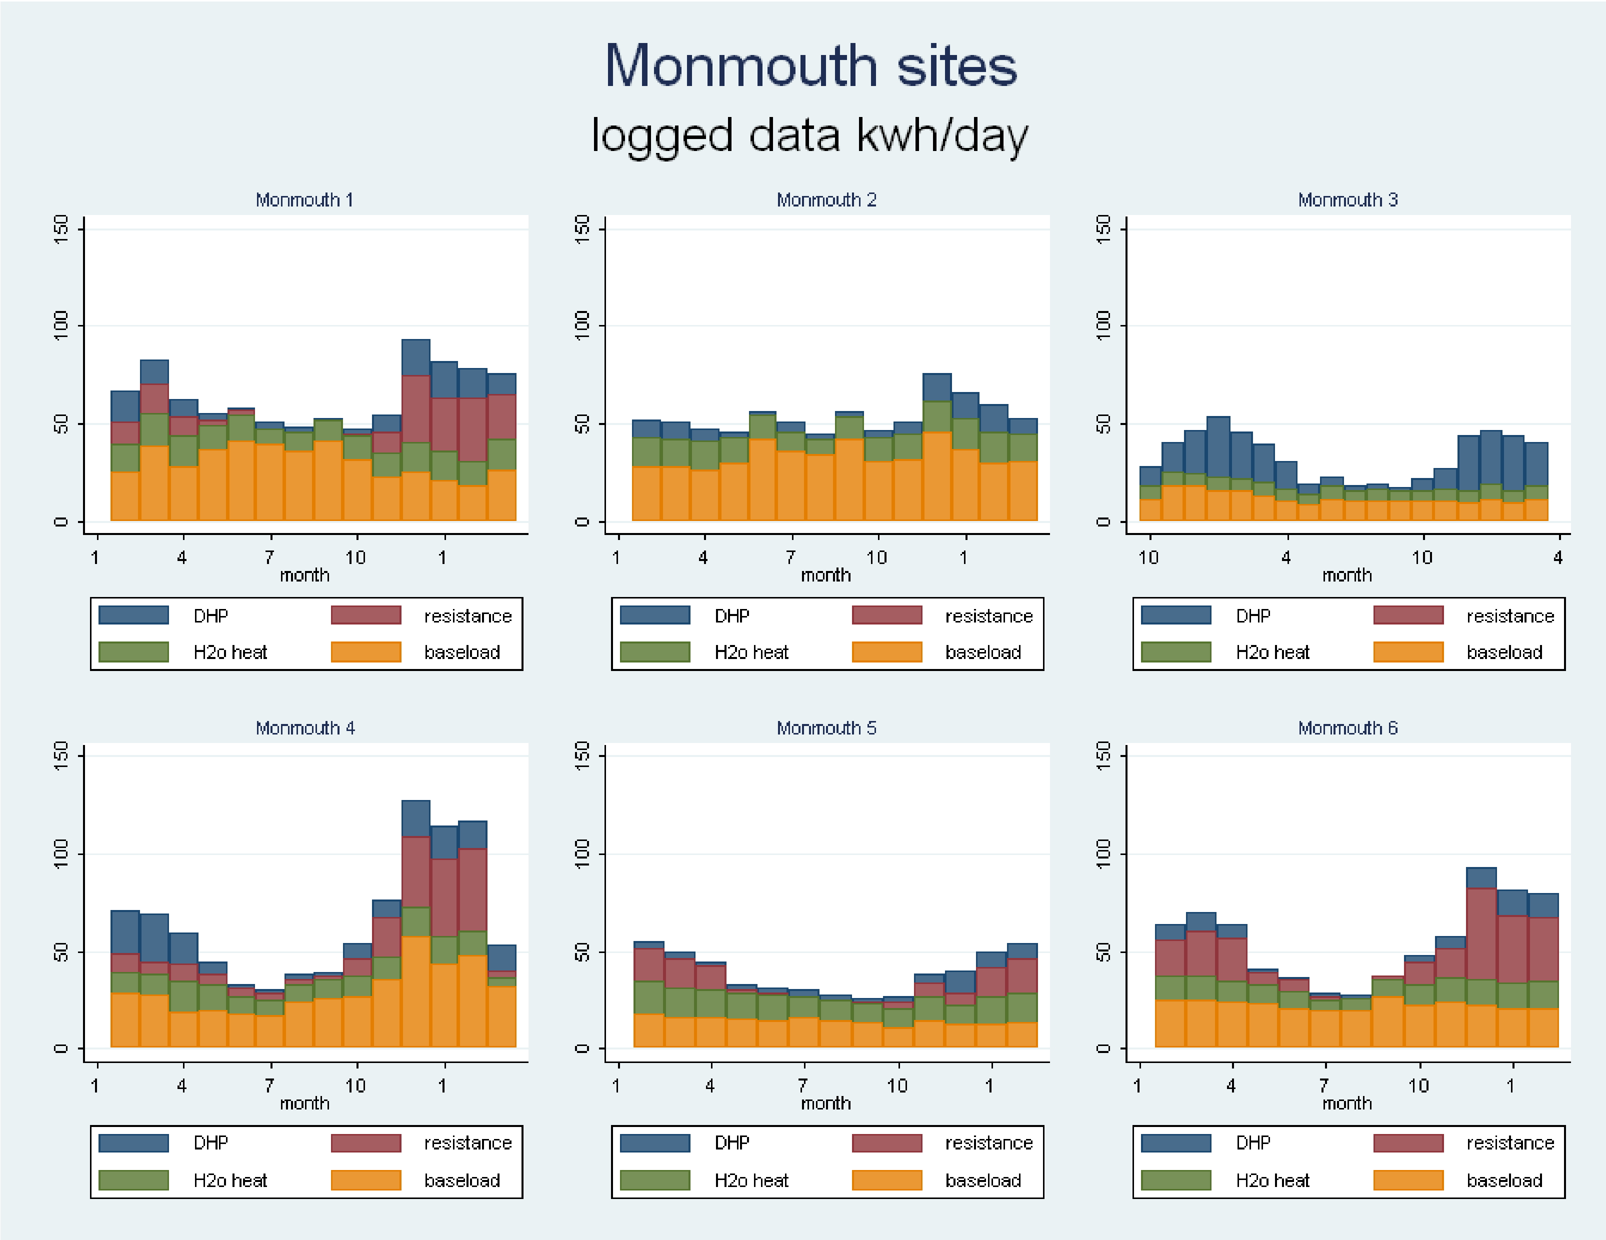
Monmouth sites
logged data kwh/day
Monmouth 1
Monmouth 2
Monmouth 3
Monmouth 4
Monmouth 5
Monmouth 6
0
50
100
150
month
DHP
resistance
H2o heat
baseload

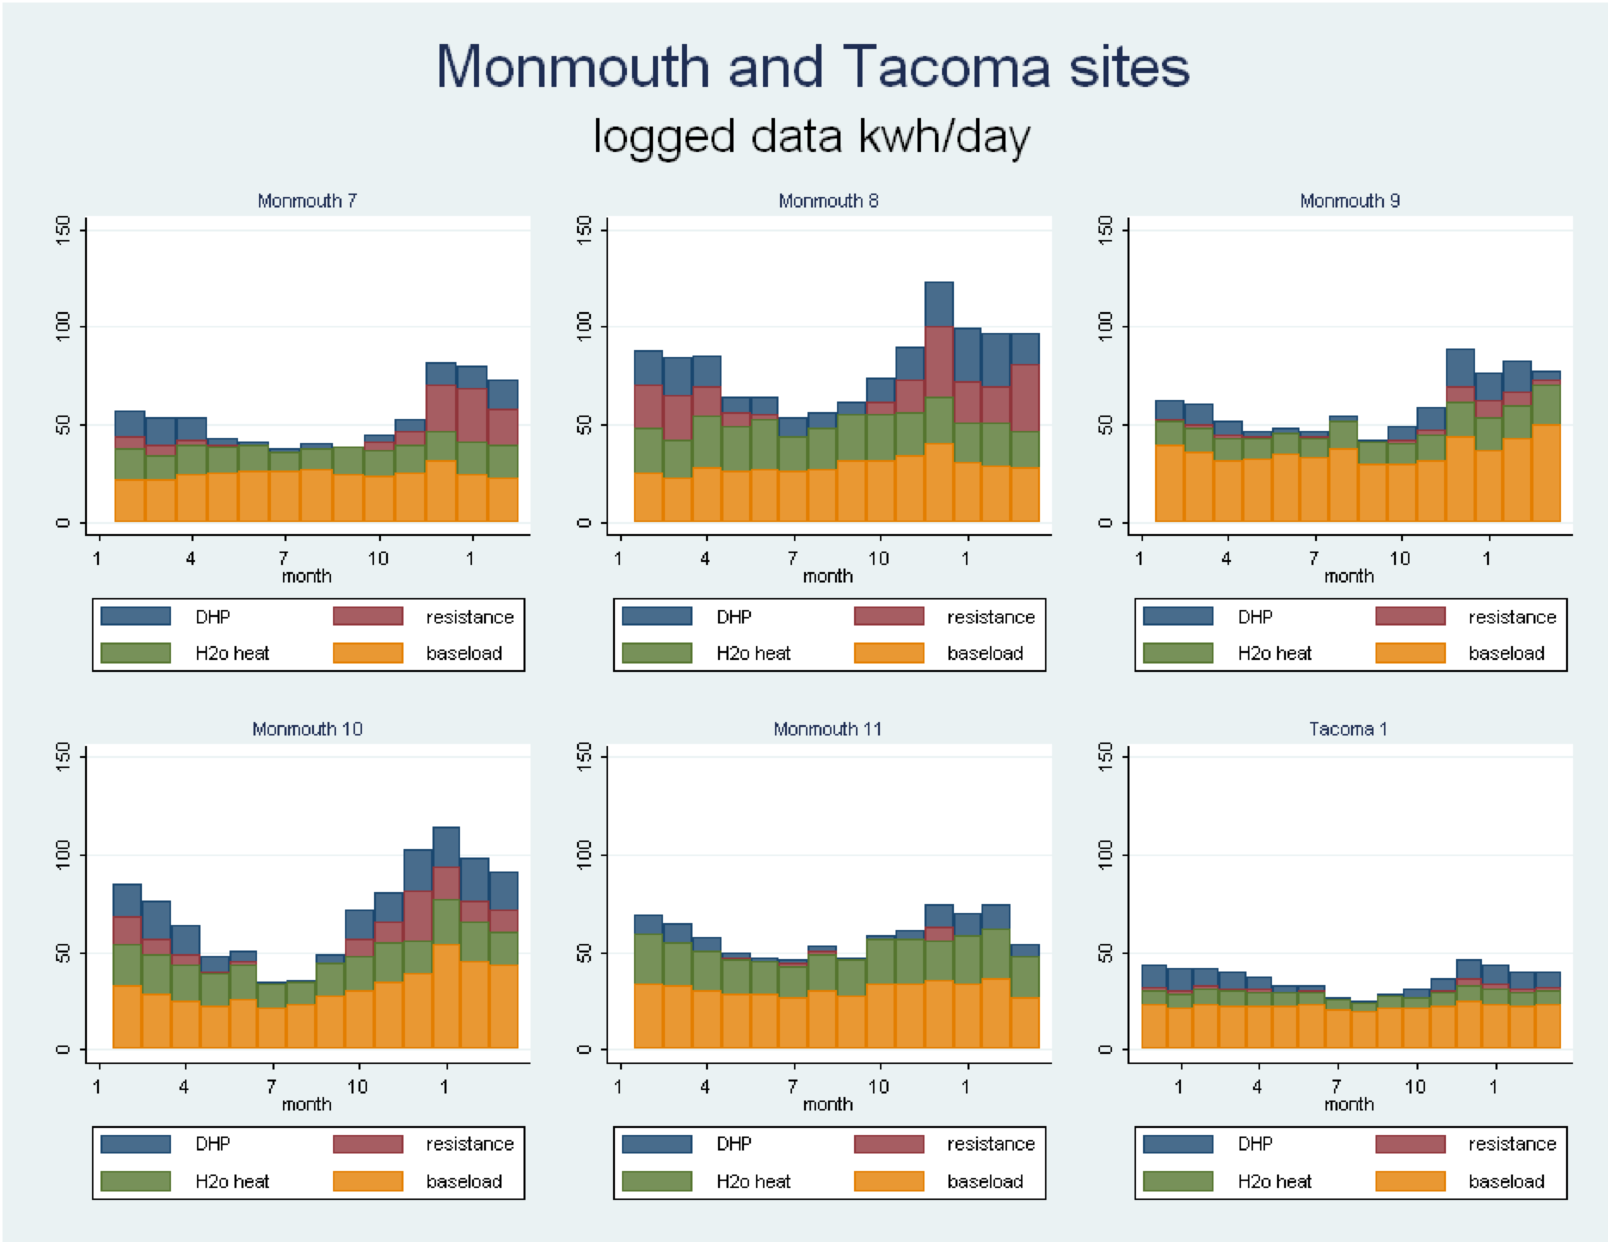
Monmouth and Tacoma sites
logged data kwh/day
Monmouth 7
Monmouth 8
Monmouth 9
Monmouth 10
Monmouth 11
Tacoma 1
0
50
100
150
1
4
7
10
month
DHP
resistance
H2o heat
baseload

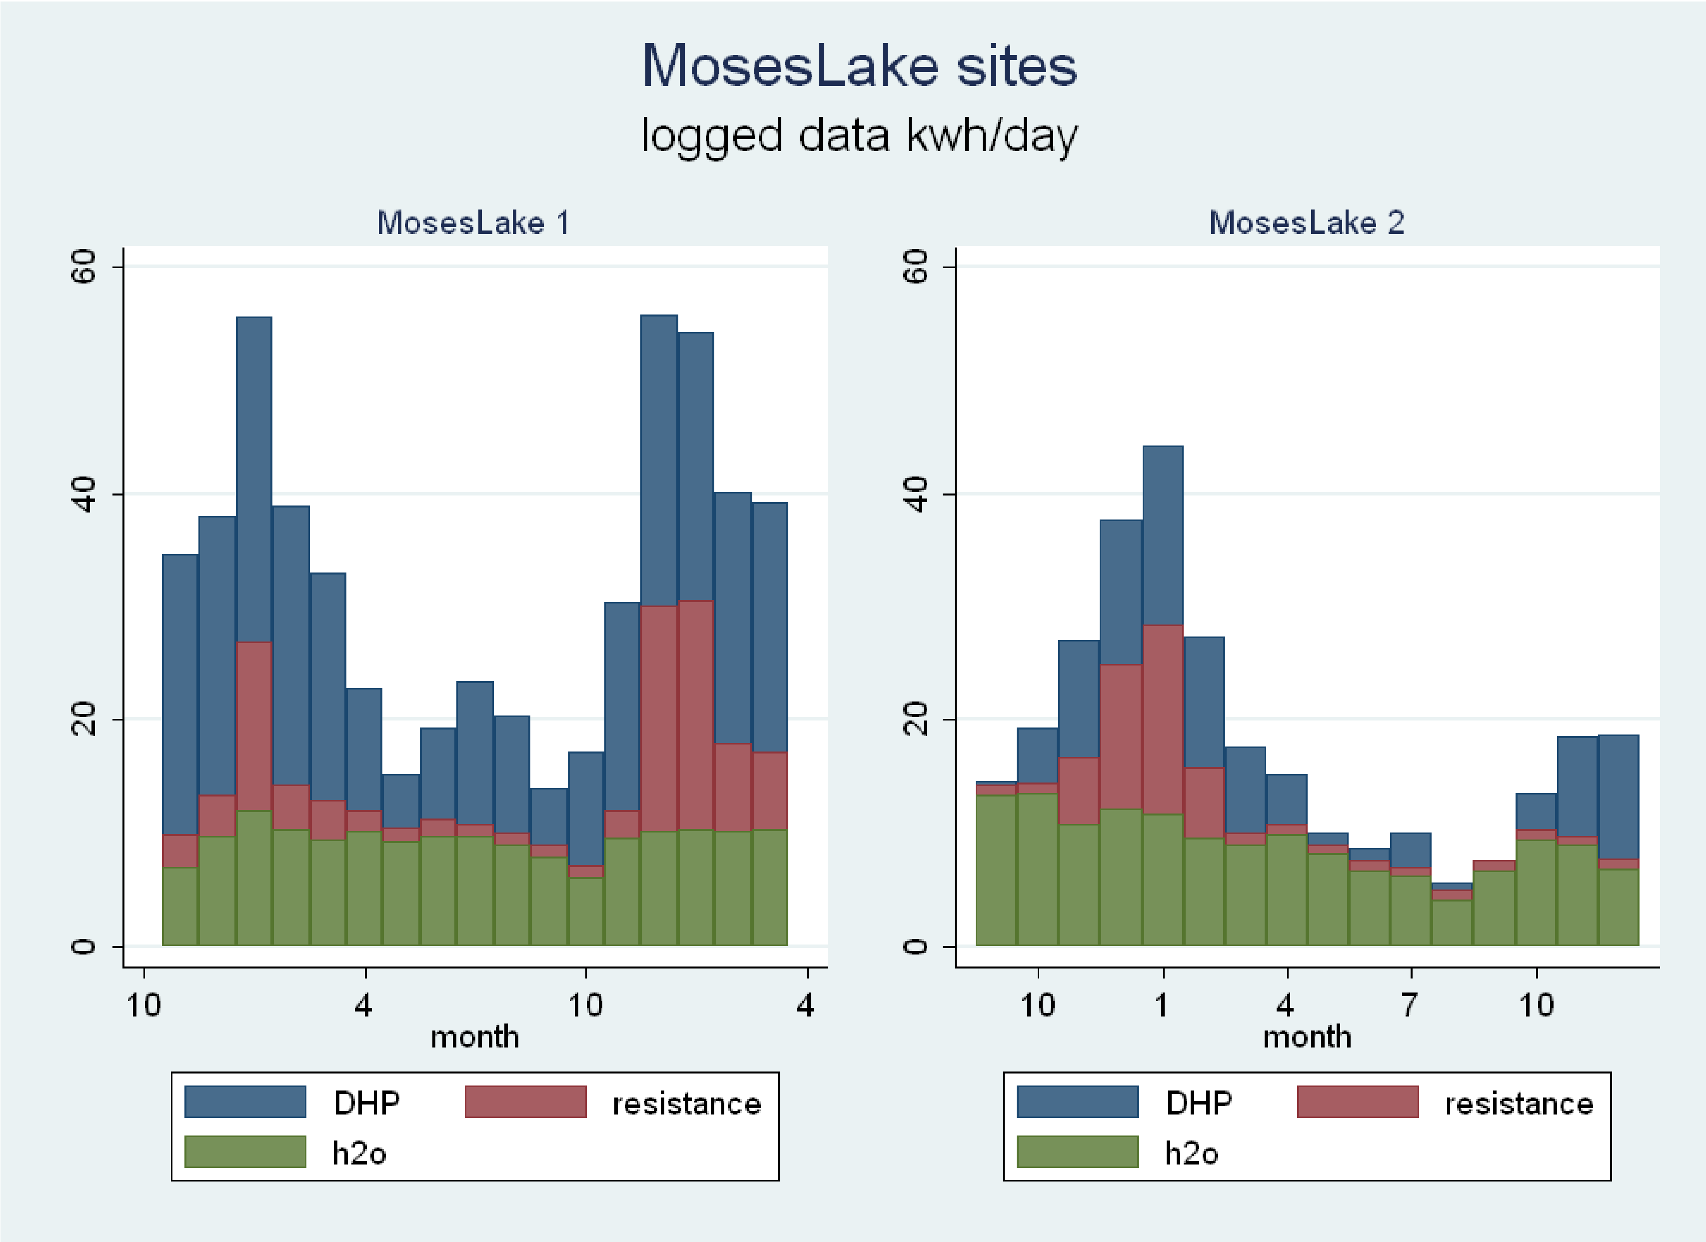
MosesLake sites
logged data kwh/day
MosesLake 1
MosesLake 2
60
40
20
0
10
4
10
4
month
10
1
4
7
10
month
DHP
resistance
h2o
DHP
resistance
h2o