Solving the Rate Problem: Fundamentals of the Metered Energy Efficiency Transaction Structure

Utility systems have long struggled with the rate impacts and core business impacts of helping customers improve the efficiency of their energy use. In an economy that increasingly values efficient energy solutions, this new approach may hold promise.

by Tom Foley and Bill Campbell

Electricity Policy – the website ElectricityPolicy.com and the newsletter Electricity Daily – together comprise an essential source of information about the forces driving change in the electric power industry.

Solving the Rate Problem: Fundamentals of the Metered Energy Efficiency Transaction Structure

Utility systems have long struggled with the rate impacts and core business impacts of helping customers improve the efficiency of their energy use. In an economy that increasingly values efficient energy solutions, this new approach may hold promise.

by Tom Foley and Bill Campbell

Utility rates are simple in concept. Regulators set rates so utilities can recover their costs, including operating costs and a fair and reasonable return of and on invested capital. [1] That means the rates utility customers pay, collectively, are basically determined by this relationship:

The utility's total costs, **divided by** all the energy units it billed, **equals** the rate per unit that the average customer pays.

If units of energy delivered (kWh or therms) go down, unless costs[2] change, revenues and rates must rise. It's simple arithmetic. [3]

Units delivered are measured at customer meters. When the utility system invests in yield improvements – either energy-producing or energy-saving devices on the utility side of the customer meter, measured customer consumption does not change. Units billed stay the same. For utility-side yield improvements, only utility costs affect rates.

Tom Foley *has over 32 years of experience in the field of energy analysis and management, including 12 years as manager of conservation and generating resources analyses for the Northwest Power Planning Council and 10 years at Battelle Northwest. He was instrumental in developing the last serious innovation in energy efficiency finance, the "public service charge" system of financing, now widely used in the United States.* **Bill Campbell** *has over 34 years of experience in energy efficiency program design, both inside and outside utilities. He is today a principal at Equilibrium Capital Group. He is an author of the Center for American Progress white paper Efficiency Works: Creating Good Jobs and New Markets through Energy Efficiency (2010) (available at* http://eq-cap.com/research/) *and holds issued and pending patents in the field.*

But yield improvements on the customer side of the meter do offset consumption, as measured at the customer meter. Measured units go down. For customer-side yield improvements, both costs borne by the utility system, and *reduced units metered,* affect rates. All else being equal, rates are higher if a resource – including an energy efficiency resource – is installed on the customer side of the meter.

This effect we call "the rate problem."[4] The "rate problem" exists whether the utility is an investor-owned "for profit" utility or a consumer-owned "not-for-profit" utility.

Designers of utility-system-sponsored programs for customer-side resources, like rooftop solar or customer efficiency, have mostly focused on the *cost* side of the rate arithmetic. Two approaches dominate.

First, utility programs try to get the customer to pick up most of the cost.[5] Second, in Oregon and other jurisdictions with a public benefits charge the portion of the cost that comes from the utility system, programs use capital does not require a return on investment.[6]

Programs supported by these cost reduction strategies have had meaningful effects. For example, Northwest utilities have acquired more than 4,600 average MW[7] (aMW) of conservation resources since 1978.[8] Over 35 years, that adds up to a large energy saving, but on an annual basis the savings may be in the hundreds of annual megawatt-hours – small by comparison to the region's total demand. This is generally true in other regions as well. Thus the "rate problem" that investments on the customer side of the meter has created is relatively small in the U.S. today.

But with a smarter grid in the offing – indeed, one is developing rapidly in some states – the distinction between loads and resources will be blurred. In a smart grid environment power will flow both ways. Many resources, some that will limit transmission and distribution needs, will be sited behind customers' meters. These resources could be generators (renewable or not) that provide power and ancillary services. Or they could be storage (for example, in electric cars), or demand response, or energy efficiency, or any other resource that provides value to the grid.[9]

Thus, the "rate problem," which has been *de minimis* to date, coupled with additional limits on opportunities for utility investment, may become a concern.

Because resources on the customer side of the meter may turn out to be significant, system-wide cost-reduction offerings alone may not suffice to keep the rate problem within manageable bounds.[10] Supplemental strategies may be required to offset rate impacts, while implementing all cost-effective resources on the customer side of the meter.

In this paper, we explore means of correcting for these rate impacts and the implied free-ridership through a new transaction system. We call it the Metered Energy Efficiency Transaction Structure.[11]

1. The Rate Problem.

The principle of physical equivalence

A kWh saved is, from the perspective of the grid, a kWh generated.[12] Energy is yielded either way. Energy saved and energy generated are interchangeable,[13] anywhere in the grid, from source to ground; there are no "special" locations.[14] But the grid is speckled with customer meters. They don't form a physical boundary. The energy in the grid may flow through them in either direction.[15] Nor do the meters materially affect physical grid operations.[16] One can connect a generator on either side of a customer meter and its output will flow into the grid just the same; one can also initiate efficiencies on either side of a delivery meter, and the grid will experience the savings just the same. The meters are only needed for delivery measurement.

The problem of economic inversion at the meter

The *economics* of the grid, however, are based on these meter measurements. Utilities use them as the basis for billing. The rate-making process uses these measurements to allocate the burden of meeting the utility's revenue requirements among customer classes for the kWh they are delivered." By convention, these meters *define* "delivered units," both for each customer and for the utility supplier.[17]

But when you install a generator or efficiency measures on the *customer* side of the meter, you are, in effect, delivering these resources to the grid *through* the meter. The problem is that its flow through the meter is backwards, with respect to the energy the supplying utility is delivering the customer at the same time. An ordinary meter can only record the net of these two flows[18] – one "positive," to the customer, and one "negative," back to the grid.[19]

So even though the resource behind the meter provides energy yield, to the customer's looks exactly the same as a reduction in use.

Because these meter measurements are used to set rates, everyone's per unit charge must then be higher to make up for the negative flow back through the meter. If the utility or a vendor has paid for or subsidized the resource, the host customer gets a windfall in the bill. And that creates the rate problem.

Understanding the rate problem

First, let's make sure we understand the arithmetic caused by the location of the meter. Then we will explore how to fix it on the unit side.

Envision a utility system that has total annual costs of $100 and that supplies 1000 units of energy per year. (Let's use kilowatt-hours, or kWh – the standard customer unit for electricity.)

When all of the units are supplied on the utility side of the meter, the rate fraction produces this result:

Table 1: Standard Delivery

Total cost of all supply:	$100.00
Total kWh *metered:*	÷ 1000 kWh
Rate per Unit:	***10¢ per kWh***
(everyone pays this rate)	

Now suppose the utility needs to move a distributed generator (say, one that generates 200 kWh).[20] It finds a new location in a building owned by "Sally," a customer, who owns a large commercial building. It leases space, moves the generator there and connects it. But it doesn't notice that the crew has inadvertently connected it on Sally's side of the meter. So the 200 kWh are flowing into the grid, but on Sally's side of the meter. Sally's use has not changed. But if the meter stays uncorrected for the connection error, here's what happens to rates:

Table 2: Delivery Inside Customer Meter, Reading Uncorrected

Total cost of all supply:	$100.00
Total kWh *metered (uncorrected):*	÷ 800 kWh
Rate per Unit:	***12.5¢ per kWh***
(everyone's rate went up)	

Remember, the utility is still delivering 1000 kWh to the grid. But it "sees" what it delivers through meters, and it can only "see" 800 kWh. So lucky Sally gets a windfall. Her cost per "measured unit" has gone up, the same as everyone else's. But the utility is charging her for 200 fewer units. All the other customers see their bills increase by 25%. Sally's rate increases by 25%, but her bill actually goes down.

This is what happens when the utility or another resource provider invests in energy yield – efficiency or generation – that gets delivered to the customer side of the meter. The energy yield gets "lost" in a way that it would not get lost if it were delivered and connected to the utility side of the meter.

The rate effect, if uncorrected, changes both rates and bills. It does not change the total cost that needs to be recovered. But it shifts the *allocation* of that total. Customers like Sally get lucky, just as a matter of where their meter is. Others see higher rates and higher bills[21].

Fixing the Problem

The rate problem arises because the delivery meter is giving the "wrong" reading. It's missing the energy that the utility – or another – delivered to or saved on Sally's side of her meter. We need to figure out how to correct the meter reading.

Solution 1: *Deliver outside the meter.*

We could run a wire from the utility's customer-site generator to the grid side of the customer's meter. Then the new resource would shows up on the "correct" side of the meter. But it may be impractical or expensive to do that.

Solution 2: *Use the same wire, and measure the two flows separately.*

A wire will deliver energy both directions at once.[22] So, as a physical matter, we can use the same wire that the customer gets energy from, to deliver energy back to the grid.

The only problem is that the meter is on this wire. The meter is not designed to separate the flows. It measures their net value. A torrent of energy could be flowing from the customer to the grid, and a similar torrent flowing to the customer, but if the two are equal the meter will measure zero.

Clearly, we need a way to distinguish between these two flows.

Distinguishing the flows. If you know how much one of the two flows is, and you know what the combined meter is reading, then you can easily determine the other flow. It doesn't matter which of the two flows you measure; you can use simple arithmetic to get the other.[23] Once you know the energy yield or energy saving of the equipment behind the meter, you can correct Sally's meter. You simply add that measurement to the meter reading to determine the delivered units. That gives you this picture:

Table 3: Delivery Inside Customer Meter, Reading Corrected

Total cost of all supply:	$100.00
Total kWh *metered (corrected)*:	÷ 1000 kWh
Rate per Unit:	***10¢ per kWh***
(no one's rate went up)	

Eureka! Table 3 looks just like Table 1. As long as we can determine what the energy yield of an installation is, we can connect the installation to the grid at the most advantageous spot, without worrying about which side of a customer meter it's on – and without distorting grid economics. All we need to do is to measure the energy yield of the installation, and if it is connected on the customer side of a customer's meter, take that measurement into account to figure customer consumption.[24]

Now let's illustrate the correction on customer bills. If we say that the customer in question, Sally, consumes 250 kWh, here's how the process just described affects customer bills, in a utility with two customers: Bob and Sally. Here's Table 1 again, the standard "grid side" calculation, with our simplified customer billing:

Table 1.1: Standard Delivery

Total cost of all supply:	$100.00
Total kWh *metered*:	÷ 1000 kWh
Rate per Unit:	10¢ per kWh
Bob: 750 kWh Consumed, 750 kWh measured ***Billing: $75.00***	**Sally:** 250 kWh consumed, 250 kWh measured ***Billing: $25.00***

Now here's Table 2 again, the "uncorrected meter" case, showing what happens to customer bills when we inadvertently supply 200 of our kWh on Sally's side of the meter, without correcting the reading:

Table 2.1: Delivery Inside Customer Meter, Reading Uncorrected

Total cost of all supply:	$100.00
Total kWh *metered*:	÷ 800 kWh
Rate per Unit:	***12.5¢ per kWh***
Bob: 750 kWh Consumed, 750 kWh measured ***Billing: $93.75***	**Sally:** 250 kWh consumed, 50 kWh measured ***Billing: $6.25***
(Generation or efficiency has been installed on Sally's side of her meter)	

Obviously Bob is annoyed and Sally is delighted. Sally gets a windfall that Bob helps pay for. Sally has not really earned this windfall.[25]

In general, whenever the system pays for (or subsidizes) installation of a resource at one customer's location, all other customers will experience an increase in their electricity bills.[26] But when we correct the meters by adding in the units delivered at the "inside the meter" delivery point, everything returns to normal.

So: as long as we can determine the energy that a particular utility resource[27] produces at its delivery point, it doesn't matter which side of a customer meter the equipment is installed on. We can just use the measurement of energy yield to correct to compensate for *where* the equipment is connected.

Here's the simple calculation:

Table 4: Delivery Inside Customer Meter

Read Sally Meter:	50 kWh
Correct for Delivery Meter:	+ 200 kWh
Sally's Corrected Consumption: (Sally's billing)	***250 kWh***

2. Measuring Energy Yield from Efficiency in General: The Principle of the Baseline Meter.

Above, we've shown that as long as you can measure the energy yield of grid resources, it does not matter where they are connected to the grid: you can use what you know about the energy yield to correct the reading on any affected delivery meter. We also showed how to make this adjustment to prevent windfalls or burdens. Making sure the energy yield is measured properly allows us to solve the rate problem. So how do we measure energy yield from efficiency?

3. Using a Meter to Measure Energy Yield in the Grid

Let's look at how to measure energy yield for "utility side" resources, whether generation or efficiency.

Direct Measurement. If the utility invests in (or buys the output of) *generation* equipment, the utility can measure the electricity from the generator directly to determine the energy yield, and to know if the generator is working.

Derived Measurement. If the utility invests in, or buys the output of, *efficiency* equipment, a direct measurement is not possible. Here the utility measures the output through a comparison of performance with the equipment and performance without it, to know if the equipment is working.

Suppose, for example, you measure the energy flows on either side of a transformer and discover it has losses of 10 units of energy out of every 100 units that enter it; the remaining 90 units flowing out the other side. The utility now installs a replacement transformer that reduces that loss to 5 units. Now the transformer uses up only 5 units instead of 10.

From the perspective of units available to sell, the new equipment has generated 5 units. But you can't install a traditional meter on the new transformer that measures those 5 units directly. The equipment's ability to create energy yield – to "generate" electricity – only works symbiotically with the transformer, and only has meaning when compared with the system it replaced. For purposes of grid performance, the yield is real. We just have to measure it indirectly. We do this by comparing how much energy the transformer previously used (here 10 units, the "baseline") ,[28] with how much it uses now.

This is the principle of the baseline meter. Here's the arithmetic:

Table 5: Measurement of Transformer Efficiency Generation Using Baseline Meter

Old (Baseline) Transformer Use:	10 kWh
Current Transformer Use:	- 5 kWh
Yield from Transformer Efficiency:	***5 kWh***

The measurement allows us to prove that the new equipment has added 5 kWh to the grid in a way every bit as real as if the utility had installed a 5 kWh generator downstream of the transformer. Our meter reveals that the efficiency "generator" is interchangeable with a generator of equal cost, operating expense, and useful life.[29] The efficiency "generator" just needs a different meter design, a "baseline meter," to determine its energy yield.

4. Applying a Baseline Meter to Measure Energy Yield to the Grid from Customer-Side Installations

What holds for the "transformer efficiency generator" is also true for efficiency anywhere in the grid, including installations that happen to be on the customer side of a meter.

To show that this is so, we take a measurement at another point on the grid – a distribution pole – and get the same answer (100 kWh in, 90 kWh out) as for the transformer. But there's no transformer here. What explains the "loss" of 10 kWh? We discover a customer line at that pole, running to a customer meter. The customer meter reads 10 kWh. Now we know: 100 kWh in, 10 kWh used, this time by the customer instead of by a transformer – so there's 90 kWh out.

But although this load is the same size as the old transformer load, and has the same physical effect on the grid, this load is behind a customer meter. That means the utility is getting paid for delivering these 10 kWh; it was not being paid for delivering the 10 kWh to the transformer. So how can we treat the customer load the same way we treated the transformer load?

We go to the customer – call her Kelly – and ask if we can harvest 5 kWh of energy on her property, and run it back up the line to the grid. She says sure, go ahead – just get my permission before you install anything, so I know it won't mess up my building.

Our goal is to have the same energy yield on the downstream side of the pole where Kelly's load attaches, as we did downstream of the pole where we installed the new transformer, so the examples are physically interchangeable. We also want to achieve economic equivalence, so the two resources are, at similar cost, also interchangeable for ratepayers. We are looking for this result:

Table 6: Desired Results, Kelly Building Energy Yield Measurements

Utility Equipment Yield:	
Old (Baseline) Use, Grid Meter:	10 kWh
New Kelly Use, Grid Meter:	- 5 kWh
Utility Equipment Yield:	***5 kWh***
Customer Meter Correction:	
Uncorrected Customer Meter:	5 kWh
Output from Utility Equipment:	**+ 5 kWh**
Kelly's Correct Consumption (Kelly's billing):	***10 kWh***

Table 6 gives us the same results as Table 5. We might go about achieving the result of Table 6 in this fashion: We will sign a lease with Kelly that lets us harvest unused energy in her building, subject to her approval for any improvements we want to make.[30] We look around for where there is unused energy and find two possible sources:

- There's sunlight on the roof that's not now being used.
- There's waste heat from many devices – light bulbs and soffit vents and single pane windows and door jambs, for example—that is serving no useful purpose.[31]

We research what equipment lets us turn these different kinds of unused energy into the 5 kWh we need in order to have 95 kWh on the other side of the pole. The research tells us we have these choices.

- A solar panel on the roof.
- WHEE (Waste heat to electric energy) units on each waste heat location.

We study both options. Both are renewable, but our (imaginary) [32] WHEE units are less expensive for the energy yield, are cheaper to maintain, have a longer useful life if maintained, and have better grid system attributes.[33] We want the WHEE units.

We also realize we can measure their actual yield contribution best, cheapest, and fairest with a whole-building baseline meter.[34] By appropriate selection of WHEE units, we can get, *exactly*, the result of Table 6 from this choice.

Then we realize we can save still more money by installing advanced WHEE units where we can – units that we can "integrate" into the heat sources themselves, so that the WHEE units can re-use the wiring the heat source already has - we save the wiring costs. The baseline meter doesn't distinguish between an integrated WHEE and a separated WHEE, so our measurement still works. Again we get, exactly, Table 6.

Finally, we realize that we can have exactly the same energy yield, and exactly the same meter readings and meter corrections, and also, because of how we have selected and designed our imaginary WHEE units exactly the same very beneficial load characteristics. By retrofitting Kelly's building we replace the equipment that generates waste heat with equipment that gives the same beneficial output more efficiently.

Our question now is whether it's cheaper to do so, or cheaper to stick with the WHEE units. Suppose it's cheaper just to replace the light bulbs with LEDs, to caulk the windows and doors, and to complete the package with new windows, and thus avoid the waste in the first place, than to install the WHEE units. This plan, too, will deliver our desired Table 6 result.

And except for the cost savings, no one will be able to distinguish between our WHEE unit solution that turns the device's waste heat back into electrons in the grid, and our efficiency solution that prevents an equal measure of electrons from being required in the first place.

So we don't install hundreds of separate WHEE units. Instead we retrofit Kelly's building with hundreds of individual measures in the form of more efficient lights, better weatherization, better windows, and so on. And we start measuring with the baseline meter. We discover that we do, indeed, have the same energy yield to the grid as in Table 5, and that, economically and physically, we have achieved the Table 6 result, at lowest cost:

Table 7: Measurement of Kelly Building Energy Yield (using efficiency generation and Baseline Meter.)

Utility Equipment Yield:	
Old (Baseline) Use, Grid Meter:	10 kWh
New Kelly Use, Grid Meter:	- 5 kWh
Utility Equipment Yield:	***5 kWh***
Customer Meter Correction:	
Uncorrected Customer Meter:	**5 kWh**
Output from Utility Equipment:	**+ 5 kWh**
Kelly's Correct Consumption (Kelly's billing):	***10 kWh***

More generally, what we have done at Kelly's building can be replicated by grid engineers, managers, and regulators in any customer facility, and will allow them to treat customer-side resources of any sort in the same way as grid resources for purposes of load analysis and rate impacts. We have also illustrated a potential solution to the rate problem.

5. Attributes of a Functioning Baseline Meter

The Bullitt Foundation's net-zero building – site of a novel baseline meter trial to prove the concept expressed here.

How, then, does a baseline meter work?

A baseline meter needs to meter the difference in consumption, in any relevant time interval at any point in the future, that is attributable to the particular changes that comprise the negawatt-hour generator. That difference is energy yield – "generation" from the perspective of the grid.

To achieve this measurement with transactional accuracy, a baseline meter must sufficiently exclude extraneous effects from so-called "neutral factors," such as changes in weather, occupancy, use, and plug load, and unrelated changes in the building. To exclude those effects, which are pervasive and often irregular, means the baseline meter system must be able to continuously recalibrate the "baseline" to adjust for the current set of neutral factors. That continuous recalibration and adjustment, with transactional accuracy, is in fact the essence of an effective baseline meter system.[35]

If a baseline meter with these characteristics is possible, the program designer can do something new: not just estimate, measure, or verify, but actually *meter* the delivered yield attributable to those changes, and so enable financing of an upgrade to the building based on that energy value.

The meter would allow the program designer to offer rigorous bill neutrality, and so get outside all the arguments about risk shifting (particularly altering mortgage holder priority) that have troubled, and blocked scale in, Commercial PACE, On Bill Recovery, and other insightful transactional finance solutions.[36]

6. The State of the Industry for Metered Energy Efficiency Transactions

Recently, Seattle City Light – the city of Seattle's municipal utility – has determined that there does exist at least one baseline meter (the DeltaMeter™ from EnergyRM[37]) that in principle can deliver such measurements, including such continuous calibration against neutral factors.[38] Based on that determination, they have entered into agreements with the Bullitt Foundation to test a 20 year "energy efficiency power purchase agreement" under the Metered Energy Efficiency Transaction Structure principles described in this paper.

The contract will be settled at the building meter of the Bullitt Foundation's new net-zero building. The settlement will rely on the baseline meter as the determinant of yield, and the readings from that meter will be used to correct the customer meter readings on the building, in the way outlined in this paper.[39]

The Seattle City Light-Bullitt Foundation contract represents a test of the meter, and a number of parties important to the Pacific Northwest's energy grid are involved in the test and in the verifying of results, primarily through the Northwest Energy Efficiency Alliance (NEEA.)[40]

7. Conclusion

In this brief and simplified paper, we have seen that grid location need not "flip the sign" on resources acquired with ratepayer dollars. Yield is yield; it is positive wherever it shows up, whether on the utility side or the customer side of a customer's meter.

Based on this principle, we have studied the impact of measurement locations on the economics of the grid. We have also noted that the meter's location can change ratepayer economics when engineers install utility installations on the customer side of a meter.

We have shown how meter readings can be used to eliminate this change by isolating the measurement of yield from those customer side resources, and, using that measurement, to adjust the customer bill so that all customer consumption is billed. We have shown that doing so in principle solves the "rate problem." And we have discussed the use and characteristics of a different meter design – a baseline meter – to provide this measurement, where the customer-side energy yield is from an efficiency resource.

We have also shown how, with this approach, we can create a transaction structure that turns energy productivity into a new grid resource: Metered Energy Efficiency.

Exploring the "rate problem" and the Metered Energy Efficiency Transaction Structure helps illuminate how an economic framework for two-way flows of energy and energy attributes can work across the customer meter. These ideas are now beginning field testing. We regard these ideas as having serious potential, best validated through transparent and well-resourced trials. If these trials produce the validation of both the measurement system and of the transaction structure being tested as their participants hope, they will earn market acceptance, the national energy grid will have a new transaction system and tool to add to its toolkit for integrating customer-side resources with grid-side resources, economically and in grid planning. □

ENDNOTES:

1 Statutes define fair and reasonable rates this way. For example, Oregon's ORS 756.040 says:

> Rates are fair and reasonable for the purposes of this subsection if the rates provide adequate revenue both for operating expenses of the public utility or telecommunications utility and for capital costs of the utility, with a return to the equity holder that is:
>
> (a) Commensurate with the return on investments in other enterprises having corresponding risks; and
>
> (b) Sufficient to ensure confidence in the financial integrity of the utility, allowing the utility to maintain its credit and attract capital.

2 Utility cost = operating costs plus return of and on capital.

3 A number of jurisdictions have adopted "decoupling" mechanisms that basically reset rates for the "lost unit" effect more often.

Decoupling lessens the burden on the utility of lost units by speeding up the rate recalculation. But it also means rate increases are felt by consumers sooner. *See, e.g.*, Arthur Rosenfeld The California Effect | commentary: Ralph Cavanagh, INNOVATIONS (MIT Press, vol. 4, no. 4, fall 2009); REVENUE REGULATION AND DECOUPLING: A GUIDE TO THEORY AND APPLICATION, Regulatory Assistance Project, June 2011. *See also* ELCON, Revenue Decoupling: A Policy Brief of the Electricity Consumers Resource Council, Jan. 2007, at http://www.elcon.org/Documents/Publications/3-1RevenueDecoupling.PDF; R.J. Michaels, Electric Revenue Decoupling Explained, Feb. 13, 2009.

[4] The "rate problem" is sometimes called the "denominator problem," because in the fraction (total cost ÷ total units,) units go in the denominator (the bottom number in a fraction.) When the bottom number in a fraction gets smaller, the answer to the fraction gets bigger. Higher rates occur whenever units are reduced if costs are not also reduced by at least a proportionate amount.

[5] Since the customer gets the benefit of the reduction in measured units (otherwise a windfall) it seems fair for the customer to pick up most of the tab. But this means customers have to decide to spend their money this way, which is a barrier to acquiring the resource. Programs generally seek to overcome this barrier through customer education -- showing the financial returns to customers, often expressed as payback periods (for example: "in just two years this improvement will pay for itself in lower energy costs"), using ratepayer funds collected via a Public Purpose Charge to pay for part of the cost of acquisition, and through subsidies for low-cost loans (for example, Clean Energy Works Oregon has arranged for 5.99% unsecured long term financing.)

[6] This strategy usually gets its capital through an up-front charge allocated by energy unit bought, to all ratepayers. The charge creates a pool ("public benefit fund") of ratepayer dollars that can be used up front; they don't have to be borrowed from someone else and paid back by ratepayers with interest over time. That eliminates cost of capital, in the same way that paying for a house up front avoids the need to pay interest on a mortgage. And as long as the public benefit charge is small by comparison to the overall bill, ratepayers tolerate it well. A notable example of such programs is the Energy Trust of Oregon, which Tom Foley helped create. The Energy Trust manages a pool of ratepayer dollars collected through a small surcharge on all investor-owned utility energy bills, to provide incentives to customers of those utilities who choose to participate, to install energy efficiency measures. Similar programs in other jurisdictions are run by the utilities themselves.

[7] In energy terms, an average megawatt is an amount of energy – in this case 4,600 aMW – times the number of hours in a year, 8,760. Thus, the Northwest's energy savings over the past 35 years in energy terms is 4,600 aMW, or 4.6×10^6 x 8,760 hours.

[8] *See* Northwest Power and Conservation Council, A New One-year Record: Regional electric energy efficiency improved dramatically in 2010) http://www.nwcouncil.org/library/report.asp?docid=641.

[9] This paper deals with the rate problem in the context of energy resources, but it is arising in the context of other grid resources as well. For example, battles over effects related to the rate problem are now in court in the US. *See, e.g.*, H. Northey and G. Nelson, *SMART GRID: Battle over 'negawatt' pricing heads to court*, H. Northey and G. Nelson, Greenwire, July 2, 2012 at http://eenews.net/public/Greenwire/2012/07/02/3, describing industry disputes over FERC's order mandating that capacity resources delivered through customer meters should be priced identically to capacity resources delivered to the grid.

[10] Cost reduction strategies that focus on eliminating cost of capital for utility resources, when coupled with loss of units in the system, can have other effects on the grid. For example, lost units reduce utility demand. In the short term, demand reductions may adversely affect a utility's return on capital, as rates may have been set based on a higher unit forecast. In the longer term, demand reduction reduces need for the capital investments that drive shareholder returns for

investor-owned utilities. Reducing return on capital, and reducing opportunities for generating shareholder return, may make it harder for utilities to attract shareholder capital and may increase a utility's cost of borrowing.

Some semi-fixes to these problems have been developed, such as decoupling, cost-recovery of lost revenue from conservation, or measures that allow some degree of utility earnings from conservation work. See, eg, *Financial Analysis of Incentive Mechanisms to Promote Energy Efficiency: Case Study of a Prototypical Southwest Utility,* P. Cappers, C. Goldman, M. Chait, G. Edgar, J. Schlegel, W. Shirley, Lawrence Berkeley National Laboratory, 2009 (http://eetd.lbl.gov/EA/EMP/reports/lbnl-1598e.pdf) and *Financial Impact of Energy Efficiency under a Federal Renewable Electricity Standard: Case Study of a Kansas "super-utility",* P. Cappers and C. Goldman, Lawrence Berkeley National Laboratory, 2009 (http://eetd.lbl.gov/EA/EMP/reports/lbnl-2924e.pdf). We are indebted to Chuck Goldman at Lawrence Berkeley National Laboratories for the apt term "semi-fix."

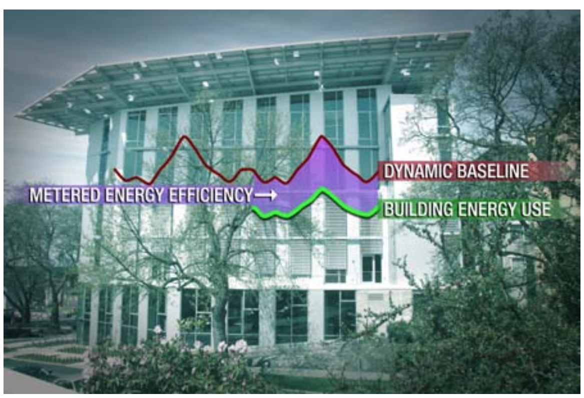

The Bullitt Foundation's living building, illustrating its use as a source of metered energy yield from energy efficiency. Photo Courtesy EnergyRM.

[11] Seattle City Light has announced the first deployment of such a system in a test of a 20 year model. Its press release is here: http://powerlines.seattle.gov/2013/06/12/new-pilot-unlocks-deep-energy-efficiency-in-buildings/.

For an exploration of how the changes explored by this paper might affect, and benefit, various parties in a customer-side energy efficiency transaction, see *Metered Energy Efficiency: A New Path to Deep Energy Retrofits,* R. Harmon and W. Campbell, 2013 (Energy Resource Management Corp. available through Harmon@en-rm.com).

[12] We describe an electric grid, but the principles apply to a natural gas grid too.

[13] Every source of energy yield has characteristic attributes which mean no two sources are ever exact matches for each other. But such attributes (load shape, intermittency, firmness, carbon load, and so on) belong to yield from generation and efficiency alike, and good utilities include both forms in their load planning and load mix, interchangeably. See, e.g., PacifiCorp 2008 IRP at 121: "Once developed, demand-side resource supply curves are treated like any other discrete supply side resource in the IRP modeling environment."

[14] There can be line losses in transmission and distribution, and bottlenecks in the transmission and distribution system, of course. All things being equal, the closer the source is to the load, the fewer line losses (leakage, for gas systems) there will be. But the physical reality of the grid takes no notice of the meter.

[15] For electricity, it's literally both directions, all the time. It's alternating current, after all.

[16] An old-fashioned electromechanical meter does draw about 2 watts of power, but this is no different than any other appliance with similar load anywhere else on the grid.

[17] Getting the "delivered units" right is critical in ratemaking. Rates are determined according to the number of units forecast for the rate period, which forecast is based on extrapolations from these measurements. Decoupling, discussed earlier, corrects for utility over- or under-collections that happen because of differences between forecasts and actual results measured at these meters.

[18] Today not all customer meters are "ordinary." One of us (Foley) has a meter that reports, separately, his consumption, and the generation of a solar panel on the roof of his house. The meter is approved by his utility, and its readings are used in aid of a net-metering transaction with that utility. But the same readings could be used to

enable a Metered Energy Efficiency transaction, as described in this paper.

[19] The term "negawatt" or "negawatt-hour" captures this picture. Of course there's nothing negative about having energy yield from efficiency. Energy yield is positive: more is more, no matter what caused it. The concept of a "negawatt" is an economic artifact created by the location of the meter, not a physical one.

[20] It doesn't matter why they need to move it, so we like to imagine it's for a nice reason, like using the old location for a new school or library or park.

[21] In the real world the effects of any one installation are much smaller than those shown in this simple example.

[22] This happens regularly on transmission lines. The ideal, most efficient, grid is one in balance: every line has energy running in both directions in equal amounts.

[23] Even if it didn't have rate consequences, you want to know the two flows separately. You need to know how much energy your generator (or efficiency) is actually delivering, to know if it is still working right.

[24] Note this measurement is not about who owns the generator. If the generator is third party owned, the utility might use this measurement to determine how many units it has received and owes for. If the generator is utility owned, the utility will use this measurement for operational monitoring. Either way the measurement will also be used to correct the reading on the customer's delivery meter.

Regulators and ratepayers will be happy as long as the "buy vs. build" choice is for the option that delivers the most bang for the buck, that is, the one that contributes the lowest total to the "total costs of supply" *ratepayers* bear.

[25] This example is about utility resources only. Sally could also install the resource on her own initiative using her own resources; she's paid for this reduction and there's no windfall. Bob's rates might still go up if she reduces her demand this way, if the utility were not able to shed equivalent costs. Then Bob would still be annoyed. We cannot solve all of Bob's problems.

[26] Utility programs that subsidize customer-side resources, like rooftop solar or energy efficiency, are generally available to all customer ratepayers. In principle all customers could take advantage of them, so "no one is hurt." (Otherwise stated – "it's their own fault if they don't take advantage of it too.") This principle has been used to justify the inequity. But it is unrealistic. There aren't proportionately equivalent and cost equivalent energy yield opportunities at every site, and even if there were, customers' ability to install them, or ability to meet the unsubsidized portion of the cost of doing so, varies dramatically. It is inevitable that some customers experience a windfall at other customers' expense. For this reason program designers add mitigating features, such as special programs for the poor.

[27] A grid resource is one the utility system is buying as energy, or has bought and paid for as capital equipment or program effect. In principle, the equipment should be owned by whoever can deliver energy so the utility's customers experience the lowest cost consistent with safe, reliable operation of the grid. (If the owner of the site wants energy from the grid, that's a separate transaction; he or she can get it at applicable tariff rates, like anyone else.)

Also in principle, any resource that the utility system helps fund with its capital – whether through a Public Benefit Fund or through purchase of energy or through any other method – is "utility equipment."

[28] Engineers will recalibrate the "old baseline" for current conditions in the period based on experience and design, to get accurate "baseline

meter" readings. In the same way, operators of a wind farm know quite accurately how energy output will vary by wind speed, air temperature, and barometric pressure, before these conditions appear.

29 Actually, it might be better. It will generate only when the transformer is working, which means energy is being used. So it's automatic load-matched generation. Any efficiency yield in principle has this load-matching effect.

30 The cost of the lease will be part of the cost of our generator.

31 Waste heat is a byproduct of inefficiently used input energy. We imagine here capturing it and turn it back into useful energy.

32 WHEE stands for Waste Heat to Electrical Energy. So far as we know, such devices do not actually exist, but they are fun to contemplate. We imagine that where waste heat exists, you can put a WHEE unit. It does not interfere with the operation of the thing producing the heat, or get in the way. It works by thermoelectric principles to turn the heat into energy. You can then just run a wire from the WHEE unit to the nearest grid-connected wire, to deliver energy to the grid. (The closest thing we found to WHEE units was a device described at www.cen-online "Turning Heat Into Electricity" Nov 14, 2011, p. 35. Alas, the described device is only about 1% efficient, so we would still be wasting 99% of the energy.)

33 Our imaginary WHEE units, because they work from waste heat, work at the same time as the thing generating the waste heat. Since the waste heat comes from use of energy, our WHEE units are automatically load-matched generation. This is useful to our utility.

34 A baseline meter is cheapest because we've spread our imaginary WHEE units all over Kelly's place -- hundreds of them, wherever there's waste heat. It would be expensive to meter them individually, or to wire them all together to a central meter. A whole building baseline meter is cheaper. It's also better, because it corrects for interactive effects. For example, waste heat from light bulbs heats the house. So in winter, lighting savings increase heating costs; in summer, lighting savings reduce cooling costs. Separately metering the light bulb WHEEs will miss the heating and cooling interactive effects. Our whole-building baseline meter will automatically net all interactive effects, revealing just the signal of the resulting yield at the pole. That's what we want to measure – the actual benefit to the system – so a baseline meter is also fairest.

35 The potential of "as delivered" energy efficiency yield was not lost on early program designers (though the concept of correcting meter readings to capture the full value of the energy harvested, discussed here, was not part of such programs). But the recalibrations were simply too problematic, so transaction systems emerged that were based on achieving essentially a delivered capacity, which could be verified, once, on delivery, and verified as "working as designed" during a warranty period. True baseline measurement was limited to particular easy-to-measure baseline cases (like individual machine replacements where load hours could be matched with high accuracy to before-and-after energy use.) For the rest, the methods of determining whether promised efficiency gains had been realized were codified in the International Performance Measurement and Verification Protocol (IPMVP).

36 PACE (Property Assessed Clean Energy) adapts the same property tax system that funds local infrastructure improvements, like sewer systems or sidewalks, to energy efficiency. Its great advantages are that the assessment transfers automatically with the property, and that the municipal lien priority reduces financing risk and therefore cost. Its great drawback is that it is not inherently "bill neutral" – the capital may or may not result in a reduction of the energy bills that offsets the increase in property taxes. That puts mortgage holders at risk, which led the Federal Housing Finance Agency to announce rules that effectively barred it in residential markets (*see http://www.fhfa.gov/webfiles/15884/PACESTMT7610.pdf*, and the companion "supervisory guidance" issued by the Office of the Comptroller of the Currency to banks: *http://www.occ.gov/news-issuances/bulletins/2010/bulletin-2010-25.html*. Commercial PACE differs because most commercial PACE programs require mortgage holder consent. Commercial PACE legal authorization exists now in 30 states and the District of Columbia; not all of these have

developed programs. (*http://pacenow.org/pace-programs/*.)

On Bill Recovery systems use the utility billing systems to collect payment for a lender for energy efficiency. To reduce lender risk and therefore financing cost, the designers push for them to be "*pari passu*" with the utility bill itself on collection (that is, any payment from the customer must be applied pro rata to the utility bill and the loan collection). If they are, they present the same "bill neutrality" problem to the mortgage industry as does PACE. If they are not *pari passu*, then they do not put building operations at risk, any more than any other unsecured loan to the building owner would. But then it is not clear whether their intended purpose of lowering financing costs will be achieved. The United Kingdom adopted an OBR system nationwide, which came into effect in early 2013; it does not appear that there has been a rush to use it. (*http://en.wikipedia.org/wiki/The_Green_Deal* and references there cited.) California has authorized an OBR program as well, which is just getting started in 2013.

37 Energy Resource Management Corp. (EnergyRM) (www.en-rm.com). One of us (Campbell) is a co-inventor of the DeltaMeter. His company (Equilibrium Capital Group) is an investor in EnergyRM and serves as its board chair. The DeltaMeter is based on EnergyRM's X-View™ Framework, the core technology of which has been widely field-tested (*http://www.en-rm.com/products-and-services/*). (The terms "X-View Framework" and "DeltaMeter", and patents both issued and pending, are properties of EnergyRM.)

38 The designers of the DeltaMeter call such a continuously recalibrated baseline a "dynamic baseline."

39 Of particular interest may be the presentation of the principles behind the DeltaMeter: T. Egnor, *The DeltaMeter™,* available at http://www.puc.state.or.us/puc/meetings/pmemos/2012/040512/Egnor.pdf (audio file also at that location.)

40 NEEA is a coalition of some 140 investor-owned and customer-owned utilities in the Pacific Northwest, along with the Energy Trust of Oregon (Oregon's public benefit fund administrator referenced in Note 7 above), and the Bonneville Power Administration.