Four Alternative Models for Regulating an Investor Owned Utility of the Future

UHERO
THE ECONOMIC RESEARCH ORGANIZATION AT THE UNIVERSITY OF HAWAIʻI

APRIL 3, 2019

How do you coerce a monopoly to act as if it were operating in a ruthlessly competitive industry? This is the billion-dollar question of Public Utilities Commissions (PUCs). It's a tricky thing to do with a mixed history of success. And it's getting trickier, especially here in Hawaiʻi, where renewable energy and so-called "distributed resources" are changing the nature of the electricity system, managed by our own monopoly, Hawaiian Electric Company.

"Distributed resources" is a catch-all phrase for rooftop solar, residential batteries, smart appliances, electric vehicles and other devices in homes and businesses that might be employed to provide power to the grid or otherwise help to keep the electricity system in perfect balance. Distributed resources convert the electric grid into a multilateral platform for trade instead of a unilateral conduit for electron delivery. The utility is still a monopoly of sorts – it still owns the wires – but it's a different kind of monopoly.

Associated with this change is growing intermittency from renewables– solar and wind–which looms large given our 100 percent renewable energy mandate. Since electricity storage is still expensive, it makes some kinds of distributed resources – flexible demand sources like electric vehicle charging – especially valuable.

There is growing recognition that these changes call for a different set of rules to govern how the utility makes its money. Ideally, these rules would encourage the utility to use its own resources when they are the least-cost way of doing things, and employ distributed resources when they are least cost. Unfortunately, under conventional rules, the utility doesn't always have an incentive to strike this balance well.

With conventional regulation, the utility essentially makes money one of two ways: (1) reducing operations and maintenance costs between rate cases; or (2) by spending it on new capital like power plants, wires, and possibly very soon, batteries.

If the PUC allows the utility to make upgrades, the utility is allowed to collect a "fair" rate of return, roughly 10% in Hawaiʻi, a rate that seems high to most impartial observers but is similar to other utilities. That is, when the utility buys capital, the PUC lets it increase rates enough to cover interest expenses plus a healthy dividend stream to shareholders. Thus, the more capital the utility buys, the higher their profit and returns to shareholders. The problem is, distributed resources and even purchase power agreements for utility-scale solar, can displace the utility's own investments. So, the utility may not employ distributed resources even when they are the least-cost way of managing the system.

How large is this disincentive? Here is a simplified example. Suppose the true cost of capital was 5 percent while the utility receives an allowance of 10 percent. Purchasing 300MW of solar through purchase power agreements would roughly displace a 75MW power plant. The capital outlay for the power plant that size would (conservatively) run about $200 million, yielding an annual profit of about $10 million per year–profit the utility loses if the solar displaces that capital. Similarly, efficiently distributed resources can displace investments in new wires, transformers, batteries or backup power plants that the utility could otherwise profit from investing in.

Many believe that efficient pricing is key to enabling distributed resources and making renewables affordable. That's basically the idea beneath a recent paper I coauthored with Imelda, a recently-minted UH Mānoa PhD and Matthias Fripp, on the faculty in Electrical Engineering Department. That paper shows that renewables can be affordable, but it requires sophisticated planning and operation, including efficient use of distributed resources, and variable marginal-cost pricing to coordinate that system.

The PUC is now considering ways to align the utility's profit motive with the broader social interests of lower costs and more renewables. To do this, it may help if we can find a way to break the link between the utility's profit and its own capital investment, and otherwise reward least-cost approaches to growing renewable energy.

There seems to be overarching consensus about these goals. There's less consensus about how to achieve them. There's also, it seems to me, some confusion on all sides.

Actually, I'm a little disappointed that my profession, which has a whole field dedicated to solving regulatory problems like this, hasn't come up with clearer guidance for PUCs. Utility regulation today looks a like Rube Goldberg device, with faint remnants of the mechanism design and agency theory I studied in graduate school. This really ought to have a reasonably simple solution. But the fact is we don't offer clear guidance; the bridge from economic theory to application remains too wide and largely ignores modern distribution-level challenges, like a multilateral grid where everyone gets to sell energy to everyone else.

Here I try summarize four ideas being kicked around for how to change the way we regulate Hawaiian Electric Company, along with some preliminary thoughts about their strengths and weaknesses.

1. Decoupling II: A broader revenue cap and less frequent rate cases.

This is an idea was forwarded by Ron Binz, former Chair of the Colorado PUC, and appears to mirror the current proposal by Hawaii's PUC staff, so hopefully I do a fair job of characterizing it. It seems roughly consistent with conventional rate-cap regulation long advocated by economists, but may require adjustment for a multilateral grid.

To explain Decoupling II, it's important to first understand Decoupling I. Currently Hawai'i, like California and a number of other states, has a rule that caps the amount of revenue that the utility can collect for operations, maintenance and some other expenses. The revenue cap replaces a rate cap–price per kilowatt hour (kWh) sold–that was the prevalent form of past regulation. People felt there was a problem with a rate cap, because it rewarded the utility for growing demand – total kWh sold. The reason utilities gain from selling more kWh is that a large share of utilities' costs are fixed, so average cost per kWh falls as total kWh rises, giving it a larger profit margin if it sells more under a rate cap.

The incentive to grow demand isn't an altogether terrible thing. To sell more and make a larger margin, it also encourages the utility to keep costs low. But it also gave the utility an incentive to discourage energy efficiency, which many believe is the least-cost way of reducing pollution emissions. So, the idea with decoupling was to break the link between the utility's profit and the kWh of electricity it sold.

It's not entirely clear how much control utilities have over energy efficiency, but we certainly wouldn't want utilities with a rate cap to be in charge of energy efficiency programs. This is the reason why a separate entity – Hawai'i Energy – manages the state's energy efficiency efforts. When rooftop solar came along, this also conflicted with the utility's incentive to grow demand for their own energy, and a key reason for revenue decoupling. Perhaps the best rationale for decoupling was the fact that demand for electricity can grow or fall for many reasons that have nothing to do with the utility's actions.

Utilities certainly didn't mind the rate cap when demand was growing, but since around 2008, demand has been shrinking, and revenue decoupling has protected the utility's profit in the face of shrinking demand. Shrinking demand and the conflict with energy efficiency and solar gave rise to Decoupling I.

But the current cap pertains only to revenue allowed for operation and maintenance of the utility. It does not pertain to revenue allowed for capital investments, which is the main source of the utility's profit. Revenue allowed for capital depends on investments actually made. As I noted above, this creates a conflict with other, potentially lower-cost ways of managing the system that involve distributed services instead of capital.

The idea with Decoupling II is to put the utility's rate of return on capital under a new larger revenue cap. No matter how much or little the utility spends on capital, they would get the same revenue. The idea is that this cap would encourage the utility to minimize capital investment and find ways to employ distributed services instead, thereby growing the utility's profit by integrating new, presumably lower-cost resources of others, like smart water heaters, appliances and electric vehicles.

This is an interesting idea. But a question some people raise concerns what happens the next time a rate case comes up and the PUC must set a new revenue cap. (A rate case, which happens about every three years in Hawaiʻi, is the legal process wherein the utility's costs are evaluated and a new revenue or rate cap is set.) The revenue cap presumably will be based on the utility's stock of capital, and what is deemed necessary going forward. There would seem to be little to prevent the utility from over-investing in the near term in anticipation of receiving a high rate of return for that investment in the future, a rate of return than many of us regard as excessive. To discourage this kind of gaming, some argue that rate cases should occur less frequently, say every five or eight years.

While less-frequent rate cases may help, potentially reducing and putting off the utility's exaggerated incentive to buy capital, I worry that this kind of revenue cap is not one that the PUC can effectively enforce. If the utility can make a case that some investments were really needed, it can legally request a rate case any time. If, between rate cases, the utility is given unfettered authority about how much to spend on capital, it will find ways to spend a lot, and then lean on an old Supreme Court ruling (and subsequent decision) that says the PUC must allow the utility an opportunity to collect revenue to procure a "fair" return on that capital.

Unnecessary capital expenses would not be hard to justify. HECO simply needs to choose not to employ distributed resources like residential batteries or demand response to relieve a congested grid, thereby creating new opportunities for HECO to spend on upgrading wires and transformers.

Also, while some of the savings from using distributed resources are capital based – like upgrading wires and transformers–some are also generation based: they allow the utility to save fuel by running their generators more efficiently. Most fuel costs, however, are a pure pass through to customers and don't fall under the cap (recently lowered to 98 percent pass through). So, to push this approach further it would need to be paired with PIMS – performance incentive mechanisms – that reward the utility for improved efficiency on other outcomes. I understand that, with a great deal of effort, current deliberations reduced the number of targeted outcomes from around 27 to "only" 12. Work is still underway to define monetizable metrics of those outcomes.

Will Decoupling II be a step in the right direction? Possibly. But to me, anyway, it seems like there needs to more to it, such as another check on capital investment besides the cap, which could complicate the Rube Goldberg device we already have. Even with this, it seems to muddle the utility's incentives at best instead of rewarding truly innovative solutions to new challenges. If we're going to have an investor-owned, profit-seeking utility, it seems we could find better ways of harnessing that profit-seeking motive for innovative, least-cost integration of renewable energy and distributed resources. I do not see how Decoupling II does that.

2. Tie a share of the utility's allowed profit to a broader, outcome-based metric.

This idea comes from my colleague Matthias Fripp. While we both feel the allowed rate of return is excessive, as do most impartial observers about which I am aware, it is also an ingrained part of utility regulation. We know of no investor-owned utility that has an allowed rate of return substantially below that of HECO's. So, if we take this high return as given, we then need to find a way to give the utility the same effective profit but does not bias their decision making toward its own capital, but instead rewards the utility for finding the least-cost way balancing the system, even if it employs others' capital and distributed resources.

Here's a simple way to do that: First, reduce the nominal rate of return on capital to something considerably lower than its current allowed rate. I think a good rate would be the average rate paid on bonds issued by investor-owned utilities. This rate is an impartial, market-based approximation for the true marginal cost of capital, and would make the utility roughly indifferent to expanding its capital base or shrinking it. It would vary over time with the larger market. By tying the cost of capital to a fixed index, it would also encourage the utility to seek the lowest cost of capital possible, for if it managed itself well, and could find a lower cost of capital, it could profit from the difference with the average utility; and it would be punished for managing itself poorly. This rate is roughly half the current allowed rate.

Second, we need to give the utility its baseline allowed revenue back, plus an opportunity to grow profit further if they can figure smart new ways to manage a high-renewable system, and punish them if they don't. Instead of tying this revenue to capital investment, it would be tied to outcomes. We might call this *shareholder revenue*. We could tie this additional revenue to a desirable outcome, like the amount of electricity used beneficially.

Here's a rough approximation: take the current allowed revenue for capital, subtract the amount allowed under the new lower allowed rate of return, and divide the remaining revenue for current capital by current kWh consumed, including that produced by rooftop solar but still connected to the grid. I believe this amount would be roughly 2 cents per kWh. Going forward, the utility's revenue cap, inclusive of both operations and capital, would include the new lower rate for capital plus 2 cents per kWh. Importantly, these 2 cents per kWh allowed revenue would not change going forward, regardless of future capital investment.

This structure bears some resemblance to the first one. It removes the direct link between the utility's profit and capital investment, for it can only collect the incremental cost of additional capital, nothing more. But unlike the first alternative, it can't try to game the system by adding more capital to grow future profit, for it cannot use capital expenditure to increase allowed revenue more than a competitively set marginal cost of capital. Instead, it obtains profit for shareholders by either providing or facilitating more of the product – electricity – regardless of how it does that.

A bit like the rate cap of old regulation, a per-kWh reward to the utility will encourage it to grow demand. But this incentive is not nearly as strong – just a small share of the average retail price. And unlike the old cap, it applies to power generated by others, since the utility must distribute, manage and balance distributed generation, not just its own sales. To grow consumption as much as possible, the utility will have a strong incentive to keep costs and prices low, regardless of who generates the power or provides grid services, for the lower the cost of electricity the more people will want to use. The utility's interests will be aligned with those providing distributed resources, so long as the distributed resources are actually less expensive than solutions that involve the utility's own capital.

A nice thing about this mechanism is that it is simple, and might alleviate some of the regulatory burden of the PUC and Consumer Advocate. It also accords nicely with the new law requiring the PUC to eliminate the explicit link between the utility's profit and its capital, and instead reward the utility based on outcomes. A small tweak – say awarding extra revenue for renewable power consumed over fossil power consumed (say 1.5 cents for fossil, 2 cents for renewable) – would be a powerful inducement for cleaner energy.

There are potential downsides to this model. First, it would discourage energy efficiency. I am not particularly concerned by this because I am skeptical that our utility has much control over energy efficiency anyway. That ball is in Hawaii Energy's court. Besides, the incentive is too low for the utility to do really perverse things like buy air conditioners for customers in order to grow demand. It might, however, be considerably more enthusiastic about developing an improved infrastructure a smart pricing for electric vehicle charging.

A second problem is that it treats all kWh of electricity the same, while in practice some kWh are much lower cost than others. As renewable penetration grows, we may have too much energy during especially sunny or windy times, and the utility may try to find ways to get people to use that energy perversely instead of curtailing (discarding it). I can imagine a utility paying people 1 cent per kWh to pump water in a circle or heating up resistors to further grow consumption and the utility's profit. They could do this and still net 1 cent per-kWh extra profit. This would be clean energy, but it's still a waste. This issue may require some policing in places where and when curtailment becomes prevalent, like it is on Maui today. Alternatively, this problem might be resolved by the third alternative.

3. Tie the utility's allowed revenue to a full-fledged estimate of net social benefits.

This is an idea I sketched out a little while ago. Policy rules always embody incentives, some desired, others not. Regardless, whatever is incentivized, will happen. So, if we want the utility to improve social outcomes connected to electricity consumption, we would be best served by tying the utility's profit to a metric that embodies a comprehensive accounting of those outcomes, preferably using the best economic measures we can muster.

This may sound like a lofty ideal; it isn't. A lot of the machinery needed to reasonably evaluate the benefits and costs of electricity use are already in place. Sophisticated generation management software is used to plan and operate electricity systems everywhere. This software needs to be revised for high-renewable systems, but that's being done too. A lot of effort has been put toward estimating the social cost of greenhouse gas emissions and other pollutants, so these can be included in the metric or priced outside of it (e.g., the polluting fuels).

The one big piece not presently common in utility regulation, but I believe critical at this juncture, is a measure of the net customer benefit from electricity use. Economists call this "consumer surplus." Consumer surplus is the most people would be willing to pay for a good or service minus the amount they actually pay. For electricity, this is probably a very big number – far larger than the actual cost of electricity, because electricity is so essential to most everything we do, yet we pay only a small share of our income for it.

Some readers may wonder what customer benefit or consumer surplus could possibly have to do with integrating renewables. What's new about customer benefit from electricity? That's a reasonable question that also has a very good answer, but it's a subtle one.

In short, tying the utility's profit to customer benefit will encourage it to employ distributed resources. More distributed resources will mean lower costs, more renewables, and happier customers. It implicitly allows for a more precise measure of marginal benefit than the per-kWh reward in option two above. This metric would vary by time of day, amount of electricity used, and other factors. Wasted electricity would have zero value, eliminating some potentially perverse incentives.

While consumer surplus is the hardest number to pin down, we can develop reasonable estimates for it, for it is embodied, to a large extent, by observed demand for electricity. In fact, we have developed a fairly sophisticated demand system and associated value metrics for electricity use in Hawai'i, a metric that can be calibrated to each hour of each day, and to the weather. (There is more demand for electricity, and more surplus, when it is hot and we enjoy relief from air conditioning.) Our demand system even attempts to account for flexible uses that can be easily rescheduled to different times of the day. We used this metric to demonstrate the potential value of variable pricing in high-renewable systems.

This calibrated demand system is not perfect. We make a lot of assumptions that cannot yet be verified. But we regard it as a reasonable starting point. What I suggest is that, even if the metric is off by a fair amount, it will still provide a good guide post for policy; it will still push incentives in approximately the right direction. Moreover, the metric can be revised over time as we learn more about behavior and work harder to link electricity use to observed weather patterns and variable prices that will be increasingly important to overall grid management.

One question about this approach is whether the utility would always be rewarded for improving consumer surplus, even if those improvements result from no explicit efforts on their own–if, for example, generation costs fall due to lower fuel or capital costs. Consumer surplus would also grow with population, since generation costs tend to fall with scale and so would the number of people enjoying the benefits from electricity use.

This perverse result would be straightforward to control using a process that is more-or-less consistent with the current regulatory process, which pins a utility's costs to model-based estimates of the least-cost way of meeting demand. The idea is to develop a modeling strategy that includes consumer surplus in addition to costs of the overall system. The rewarded metric would be a share of the difference between what is achieved and what the model estimates to be optimal. Thus, if costs fall or population grows, the model will adjust the baseline accordingly.

In a way the utility is already rewarded for a similar metric. Under revenue decoupling it gets to keep cost reductions from the base year, regardless of what causes those cost reductions, like changes in scale or technology, usually excluding capital costs. What I am proposing here is 1) a more comprehensive model that includes customer benefit; 2) a baseline that adjusts over time with technology and the evolving economy; 3) a reward for improvement that need not be dollar-for-dollar, such that the public and the utility share the value of improvements.

The main weaknesses of this approach are that precise measures of consumer surplus do not exist and it is untested as a performance metric. It would be worthwhile to simulate the approach in more detail to see how large these weaknesses might be. For example, some research could examine how distorted incentives would be if we used an incorrect metric of consumer surplus. My hunch is that this won't matter too much, but a considerable research effort is needed to test this hunch.

4. A regulated, cost-based market.

When I discuss the topic of utility regulation and incentives with colleagues more experienced in energy economics than I am (most of my research is in agricultural economics), I generally encounter reticence. They all realize how complicated the regulatory process is and the perverse incentives of rate-of-return regulation. Why don't more economists push for regulation that gives utilities better incentives?

My sense after talking with many of them is that they fear that incentive design is tricky and inevitably leads to "can't lose" situation for utilities. A poor calibration of incentives could lead to a huge windfall for utilities and a loss for customers, a situation that could be difficult to reverse. I suspect this risk is greatest when using the first approach, or what I earlier made fun of as "pocket full of PIMS". And if the utility does poorly, the utility can sue to get their fair return anyway, using the Supreme Court precedent linked above.

In other words, heads the utility wins, tails the customer loses.

Instead, the best energy economists I know suggest that we create markets when sufficient competition is viable and regulate with careful oversight when they are not. Market design is another complex topic, one that needs reconsideration with renewables and storage. Paul Joskow, one of the nation's foremost energy economists, has recent paper describing these issues. But they are largely irrelevant to Hawaiʻi because we are surely too small to have a viable competitive market.

Frank Wolak, another leading energy economist who visited Hawaiʻi last Fall, argues that places like Hawaiʻi ought to have regulated "cost based" markets. You can find a copy of his slides and a video of his talk here. The idea basically creates a competitive forward market for electricity, far enough into the future that new generation can be built. This should bring about a large degree of competition, even in a small market like Hawaiʻi, since anyone can enter and participate. This works in a similar fashion to a purchase power agreement, but Frank argues that we do not need all of the details of specific projects. Participants simply need to guarantee delivery of a proposed amount of electricity at a proposed price (a certain share of renewables can be required). Then, in real time, a computer program takes offers from any and all who want to participate, including real-time offers together with forward commitments, selecting the least cost sources that collectively meet any renewable portfolio standard. The incremental price that clears demand in this managed, cost-based market is given to distributed resources, such as homeowners who feed solar into the grid or strategically discharge batteries.

There are a number of examples of cost-based markets that work in this manner, such as Chile, Bolivia, Peru and most Central American countries. But it's not clear that any use these markets price distributed resources in the manner that we suggest here.

Presumably we would still have cost-of-service or decoupling regulation for operations, maintenance, wire, transformer and other grid upgrades, much as we currently do. But now all distributed resources would be rewarded appropriately, so it should be clearer where and when grid upgrades are actually required. The utility could also compete in forward markets for power generation, and earn a return on those investments, but only if their offers were competitive in the forward market. No regulated rate of return needs to be applied.

One wrinkle here, one that is of keen interest to Frank Wolak and other top energy economists, is how to adjust market clearing prices in different parts of the connected grid, especially when some circuits have congestion (in-flow or out-flow approaching wire capacities) or voltage regulation challenges (a problem that can arise in circuits with a lot of solar photovoltaics). It will be important to adjust prices for distributed resources in a manner that takes these issues into account. Without those adjustments, the utility might still engage in unnecessary upgrades. But this is probably a relatively minor issue if system-wide prices are set well in real time using a cost-based market.

Some might also worry about having sufficient participation and sufficient commitment from those participating in a forward market, but these are issues that can and have been managed effectively.

Which alternative is best, and where do we go from here?

Frankly, I am not quite sure. But I have ordered these four proposals in reverse order of my current preference. This order also happens be in reverse order of the size of the institutional changes that would be required, and reverse order of my perceived likelihood of being adopted.

Change is hard. We have done a lot. But we still have a long way to go.

Michael J. Roberts
UHERO Research Fellow and Professor of Economics