Algorithmically developing efficient time-of-use electricity rates

Benjamin Whitney Griffiths | 11-18-2020

Abstract (75 Words)

Time-of-use ("TOU") electricity rates are challenging to design and results often rely on a regulator's subjective sense of reasonableness. Here, I describe an objective method to develop TOU rates using optimization techniques. The method identifies rates which minimize the difference between TOU and hourly costs, thereby ensuring rates are efficient and cost based. The technique also enables direct comparison of competing designs. I conclude by demonstrating how this methodology could be applied in practice.

Abstract (200 Words)

The disconnect between price variability of wholesale electricity sales and the uniformity of the retail prices observed by most consumers results in well documented inefficiencies. Time-of-use ("TOU") electricity tariffs offer a middle-ground between dynamic wholesale prices and flat retail rates. A customer subject to a TOU rate has their consumption binned into discrete seasons (e.g., summer, winter) and periods (e.g., off-peak, mid, on-peak) – with prices varying by season and period. TOU rates can be challenging to design and the results often rely on a regulator's subjective sense of reasonableness.

In this paper, I outline a new method for developing multi-season, multi-period TOU rates using optimization techniques. This method identifies optimal TOU tariffs based on granular price and load data, using a curve-fitting approach. These TOU tariffs include explicit tradeoffs between different cost categories and between different season or period specifications. It also offers a simple method for comparing the relative efficiency of alternative rate designs, even if those designs have different numbers of seasons, periods, or price differentials. After describing the approach, I apply this methodology to a utility in California to demonstrate how it *could* be implemented in practice.

Keywords

Rate Design; time-of-use rates; Tariffs; Optimization

1. Introduction

Prices change rapidly in wholesale electricity markets, as the balance of supply and demand constantly changes. The real-time prices in competitive markets can vary dramatically from month-to-month and, increasingly, from minute-to-minute. By contrast, the prices exposed to most end-user customers are generally flat for months or years. Traditional electricity rates have customers paying the same volumetric rate (cents/kWh) irrespective of *when* that electricity was consumed over the course of a billing period. These flat volumetric rates may vary by season. Depending on your perspective, these flat rates are beneficial (because they are simple to understand and protect customers from excess price volatility) or harmful (because they abstract away price signals that could encourage more efficient use of electricity). Continued reliance on flat rates is also a product of legacy metering or billing systems which cannot handle multiple rate periods.

Irrespective of the rationale, the disconnect between the uniformity of retail sales and the variability of wholesale sales results in well documented inefficiencies.[i,ii] Consumers on flat retail rates are not price-sensitive (beyond their total monthly bill) so have no incentive to curtail their load as prices rise or system reliability is threatened.[iii,iv,v,vi] Consumers do not have incentives to shift their discretionary or flexible energy consumption to lower cost periods which would reduce the overall cost of producing electricity.[vii] Flat rates may be more expensive for end-user customers over the long run, than time-varying rates due to risk premia, hedging, and other costs.[viii] Flat rates may lead to suboptimal buildout of new supply resources[ix] or inhibit the deployment of otherwise productive uses of low-cost energy (e.g. electric vehicles).[x,xi] And so on.

Time-of-use ("TOU") electricity tariffs offer a middle-ground between dynamic wholesale prices and flat retail rates. A customer subject to a TOU rate has their consumption binned into discrete seasons (summer / winter) and periods (e.g., off-peak, mid, on-peak) – with prices varying by season and period. On-peak periods will be more expensive than off-peak; and one season will typically be more expensive than another.

The TOU concept was developed more than 50 years ago and has been rolled out to larger commercial and industrial customers in many regions.[xii] TOU rates remain rare, however, on a residential or small-commercial level.[xiii] While TOU rates are less efficient than real-time pricing, they can still address many of the identified inefficiencies of flat rates. They can increase efficiency and improve price formation, reduce peak demands, and promote mid-day load growth to counteract the solar-induced "duck curve."[xiv] TOU rates rolled out in pilot programs or across utility service territories have been shown to effectively shift power consumption away from peak periods – saving money for customers, utilities, and society.[xv]

Designing TOU rates can be challenging, however. Into how many seasons and periods should the year be divided? More granular time-periods may better reflect the underlying real-time prices but are also harder for consumers to understand. After settling on the overall number of seasons/periods, what months should be in which season, and which hours in which period? Various methods have been developed to assess the efficacy of a TOU rate but these do not always provide insight how to design a rate in the first place.[xvi,xvii] For those more fundamental questions, we have long-standing guides such as Bonbright[xviii] or more recent frameworks and best-practices from RAP[xix,xx], RMI[xxi], Synapse[xxii], or others. The range of design considerations, evaluation metrics, and embedded policy goals, lead regulators to adopt "close-enough" rates – those where the periods look reasonable within some subjective range.[xxiii]

In this paper, I outline a new, objective method for developing multi-season, multi-period time-of-use rates. This method algorithmically identifies fixed periods and seasons which most closely reflect the underlying, noisy hourly wholesale market energy prices. It also offers a simple and objective method for comparing the relative efficiency of different candidate rate designs. After describing the approach, I apply this methodology to generic utility in California to demonstrate how it *could* be implemented in practice.

2. Theory & Calculation for Establishing TOU Seasons, Periods, & Prices

This TOU estimation technique identifies the seasons, periods, and prices which, when considered as a whole, minimizes the price deviations between the underlying hourly prices and the simplified TOU rate schedule.

It searches across the set of possible TOU specifications adjusting the periods and seasons of our TOU schedule in pursuit of a solution which minimizes the Sum Squared Error ("SSE") metric between hourly LMPs and the TOU prices. Due to possible non-convexities in the search space, the algorithm relies on a random-restart hill-climbing search technique which finds local minima based on many trials with randomly selected starting parameterizations.

2.1 Pricing a Specific TOU Specification

The algorithm splits the hours of the year into discrete seasons and period. The algorithm starts by cutting the year into S seasons (here, $S = 2$). More specifically, the algorithm selects the start-month of each of the S seasons.

$$\overbrace{m_a \quad \cdots \quad m_{b-1}}^{Season\ A} \overset{\longleftrightarrow}{|} \overbrace{m_b \quad \cdots \quad m_{a-1}}^{Season\ B}$$

Assuming the cuts occur in February and June; then the year is split into two seasons with Season A running February though May, and Season B running June through January. If the second cut was in July instead of June, then the Season A would be a month longer and Season B a month shorter. Additional "cuts" could be incorporated to create a tariff with more than two seasons.

Next, the algorithm splits the hours of each season of the day into P periods (here, $P = 3$). As before, the hours of the day are cut, with each cut reflecting the first hour of a given period. More periods can be added by inserting more "cuts".

$$\overbrace{h_a \quad \cdots \quad h_{b-1}}^{Period\ A} \overset{\longleftrightarrow}{|} \overbrace{h_b \quad \cdots \quad h_{c-1}}^{Period\ B} \overset{\longleftrightarrow}{|} \overbrace{h_c \quad \cdots \quad h_{a-1}}^{Period\ C}$$

Each season can have a different period specification and, indeed, different number of periods. Assuming that the hours of the year are cut at hours 3, 12, and 16, we are left with three periods running (Hours 3-11; 12-15; 16-24 and 1-2). Note, at present, the algorithm does not make a distinction between weekdays and weekends when segmenting periods, although this could be incorporated if warranted.

Using this general approach, all hours of the year can be grouped into a set of contiguous seasons and periods where are both mutually exclusive and collectively exhaustive.

Having identified each hour as belonging to a specific season and a specific period, the algorithm calculates the price of delivering energy in that season/period. Energy prices and customer demand varies hour-to-hour, but the TOU rate must simplify that variability into a much smaller set rate schedule. The TOU period's rate equals the real-time load-weighted price. Each season/period has a distinct price based on the underlying fluctuation in market prices and demand. Importantly, the total cost to serve load is the same under the real-time price and the TOU rate.[1] Or, put more simply, over the course of the year, the TOU rate generates total revenue equal to the real-time cost of serving customer demand (assuming that existing load is not price responsive).[2]

2.2 Assessing the Goodness-of-Fit of a TOU Specification

The period prices developed in the preceding section are averages. While they accurately reflect periods *average* price of electricity, there may be substantial deviations between the real-time price and the TOU rate on an hour-to-hour basis. Imagine, for example, a period where real-time prices all equally likely to be \$20/MWh or \$100/MWh. The average price will be \$60/MWh – but in half of hours the TOU price will be \$40/MWh too high and in half of hours it will be \$40/MWh too low. While costs would be covered on average, the average price does a poor job of reflecting the variability in the underlying real-time prices.

All else equal, a better TOU tariff should more closely mirror real-time prices and have deviations of lower magnitude. I capture the magnitude of deviations using the common Sum Squared Error ("SSE") metric:

$$SSE = \sum_{h=0}^{H} (Price_{h,realtime} - Price_{h,TOU})^2$$

SSE measures how much the TOU varies from the real-time prices across the full study period, putting added weight on the periods when the two prices diverge more. If the real-time prices closely mirror the TOU price, then the SSE will be small. If real-time prices fluctuate widely around the TOU price, however, residuals and SSE will be high. On a relative basis, a TOU rate which more closely mirrors the underlying real-time prices will have a lower SSE and be deemed a better fit.

2.3 Searching for High Quality TOU Specifications using Random-Restart Hill-Climbing

The preceding sections provide a framework to evaluate a specific TOU specification, but another tool is needed to search for TOU parameterizations which closely fit the underlying dataset. There

---

[1] $\sum_{h=0}^{H}\left(Price_{h,realtime} \times Demand_h\right) = \sum_{h=0}^{H}\left(Price_{h,TOU} \times Demand_h\right)$

[2] Extending this model to account for consumer response to the TOU is a conceptually a simple addition.

are many possible TOU specifications – enough that it is infeasible to assess all combinations directly.[3]

Instead of brute-force trial and error, it is faster and easier to use a search algorithm to identify high quality specifications. In this implementation, I rely on a on a random-restart hill-climbing algorithm which finds local minima based on many trials with randomly selected starting parameterizations.[xxiv] The algorithm initializes a trial with a randomly selected TOU specification and assesses the SSE of this starting parameterization. The steepest-ascent hill climbing algorithm searches all adjacent TOU specifications (parameter $\pm$ 1 hour or $\pm$1 month) and, as possible, the moves to the adjacent specification with the lowest SSE. It will repeat this search process – updating the specification along the way – until no further improvements can be found. After many trials, each with many iterative steps and even more adjacent assessments, the TOU specification with the lowest SSE is deemed optimal. (If there are nonconvexities in the search space, local-search routines like this hill-climber can get stuck in local minima rather than finding the true the global minimum; repeated trials with random initialization help ensure that the global minimum is identified.)

Because the search algorithm endogenizing tradeoffs across seasons and periods, it can generate optimal TOU specifications which closely reflect the underlying real-time prices and underlying marginal costs.

2.4 Incorporating Economic or Policy Constraints

Evaluations of TOU efficacy have shown that consumers respond better to relatively short peak periods and to higher differentials between off-peak and peak period prices. [xxv][xxvi] While the TOU rates modeled here are designed to reflect marginal costs so that as customers respond to a price signal aligned with actual system costs, policy constraints can be added to the search routine. These could take many forms. For example, minimum or maximum period lengths can be required to ensure that a peak period fits into some range of possible durations. Shoulder months could be forced into certain season. TOUs where the peak/off-peak price ratio is close to 1:1 could be discarded, because they are unlikely to shift consumption or reduce peak demands.

2.5 Comparing Results from different TOU Specifications

Using the techniques developed above, it is possible to identify an efficient, cost-based TOU for a given TOU specification. But what if we want to compare how the best 2-season/2-period specification compares to the best 2-season/3-period alternative? How should we decide which specification is better and how can we tell? All else equal, increasing the number of TOU periods and seasons should decrease SSE values, implying improved goodness of fit. At the extreme, a TOU which has the same number of periods/seasons as it has hours will have an SSE of zero (because each TOU period has its price equal to the hourly real-time price). As with traditional statistical models, there is the possibility of "overfitting" the TOU. This has two problems.

---

[3] For a 2-Seasons / 3-Period TOU, there are more than 270 million possible specifications: $\left(\frac{24!}{3!(24-3)!}\right)^2 \times \left(\frac{12!}{2!(12-2)!}\right)$. The problem gets much harder still, with larger numbers of seasons or periods.

On a substantive level, a highly granular TOU with many seasons/periods may be less understandable (or actionable) to customers, violating one of Bonbright's canonical criteria of desirable rate structures.[xxvii] If there are too many periods, regular customers may find it difficult to effectively schedule or defer energy intensive tasks to lower-price periods. A simpler tariff may provide more actionable price signals – at the expensive of some efficiency. On a technical level, if the TOU if "overfit" such that it closely mirrors the underlying historic ("training") dataset, then the TOU specification may fare poorly when implemented. Overfitting a TOU could result in the utility significantly over-collecting or under-collecting its revenue requirement requiring more active stewardship.[xxviii]

To overcome these challenges and important qualitative considerations, it is possible to adapt standard statistical tools to identify how differing model specifications fare compared to one-another. Relative model quality can be assessed using metrics such as the Akaike information criterion (AIC) or the Bayesian information criterion (BIC). For a set of candidate specifications, AIC and BIC estimate the quality of each model, relative to each of the other models. Each of these metrics includes a term which penalizes "bigger" models (here, those TOU specifications with more seasons and periods) and allow for explicit tradeoffs between model specificity and parsimony. In both cases, a smaller AIC or BIC value indicates a "better" model. BIC will penalize a "bigger" model more than AIC. The general formulas for AIC and BIC are:

$$\text{AIC} = n \times \ln(SSE/n) + 2 \times k$$

$$\text{BIC} = n \times \ln(SSE/n) + k \times \ln(n)$$

Where,

SSE is the model's sum squared error,
n is the number of values in the dataset, and
k is the number of parameters in the TOU specification.

In practice, a TOU specification is a kind of model specification, so AIC and BIC can be used to directly compare the quality of different TOU specifications and allow for the selection of the most parsimonious but information rich TOU. It is not the absolute size of the AIC/BIC value that is important, but the relative values – and relative differences -- over the set of models considered. Thus, we can use AIC/BIC metrics for direct comparison of competing rate designs, even if those designs have different seasons, periods, or price differentials. The "better" design is the one which has a lower AIC or BIC metric than its competitor.

2.6 Mostly Endogenous TOU Selection

Section 2.3 offers a method to find high-quality TOU specifications that correspond to exogenous assumptions about the number of seasons and periods. This constrained search is helpful if installed electricity meters have constraints on the number of periods/seasons it can track.[4] If metering does not pose a technical challenge, however, exogenous requirements on TOU specifications constraint can be easily relaxed. In the same way that Section 2.3 randomly initialize

[4] For example, some early interval meters can only track four distinct time periods – and reprograming them on the fly can pose technical challenges.

trial parameters for a given TOU specification, it is also possible can also randomly initialize the TOU specification itself.

Running random specifications through the search algorithm, then ranking results based on AIC/BIC metrics, provides a method to even more generally search across the whole domain of possible TOU rates. In this most basic approach, rates can be established using four sets of constraints:

- minimum and maximum number of seasons;
- minimum and maximum length of each season;
- minimum and maximum number of periods within a season; and,
- minimum and maximum length of each period.

For example, if it is decided that a TOU rate must have 1-3 seasons and that each season must have 1-3 periods, then it is possible to assess a range of different TOU specifications ranging from a single year-round price up to a TOU with 9 discrete time/price segments.'

3. Calculation & Discussion for a Sample a California Utility

Section 2 outlines a method to identify efficient TOU rates. Now, I show how it could be implemented in practice. Section 3.1 describes the dataset upon which the TOU search algorithm is run. Section 3.2 demonstrates how to apply the technique when developing a TOU with a known number of seasons and periods (cf. Section 2.3). Section 3.3 demonstrates how to compare the optimal rates from two different TOU specifications (cf. Section 2.5). Section 3.4 demonstrates how the technique could be applied using the blank slate approach (cf. Section 2.6). Please note that my purpose is not to suggest that the identified tariffs are "correct," rather it is to demonstrate how the general methodology *could* be implemented in practice.

3.1 Shape of Data

The TOU search algorithm requires costing data on whichever functions are shaped within the TOU. If transmission and distribution costs are included in the TOU, then these functions much be shaped through some sort of cost allocation process which provides hourly T&D costs, ideally on a class-level.[xxix] For this analysis, I pair hourly load profiles with prospective marginal cost data from a California utility used in a rate-design docket.[xxx] This utility resolved many nettlesome T&D and capacity cost allocation issues, and developed a dataset of hourly marginal costs for energy, capacity, and distribution. Note that the specific prices and loads relied on in the section are less important than the concepts behind them. Different cost allocations or load profiles may, of course, lead to materially different outcomes.

Figure 1 plots the costs, by category, over the 2018-2019 timeframe. While energy costs generally stay in a relatively tight band around $31/MWh, capacity costs are very high in a few hours (up to $20,000/MWh) but otherwise zero. Distribution costs fall somewhere in between.

Figure 1: Hourly Prices by Category, 2018-2019 ($/MWh)

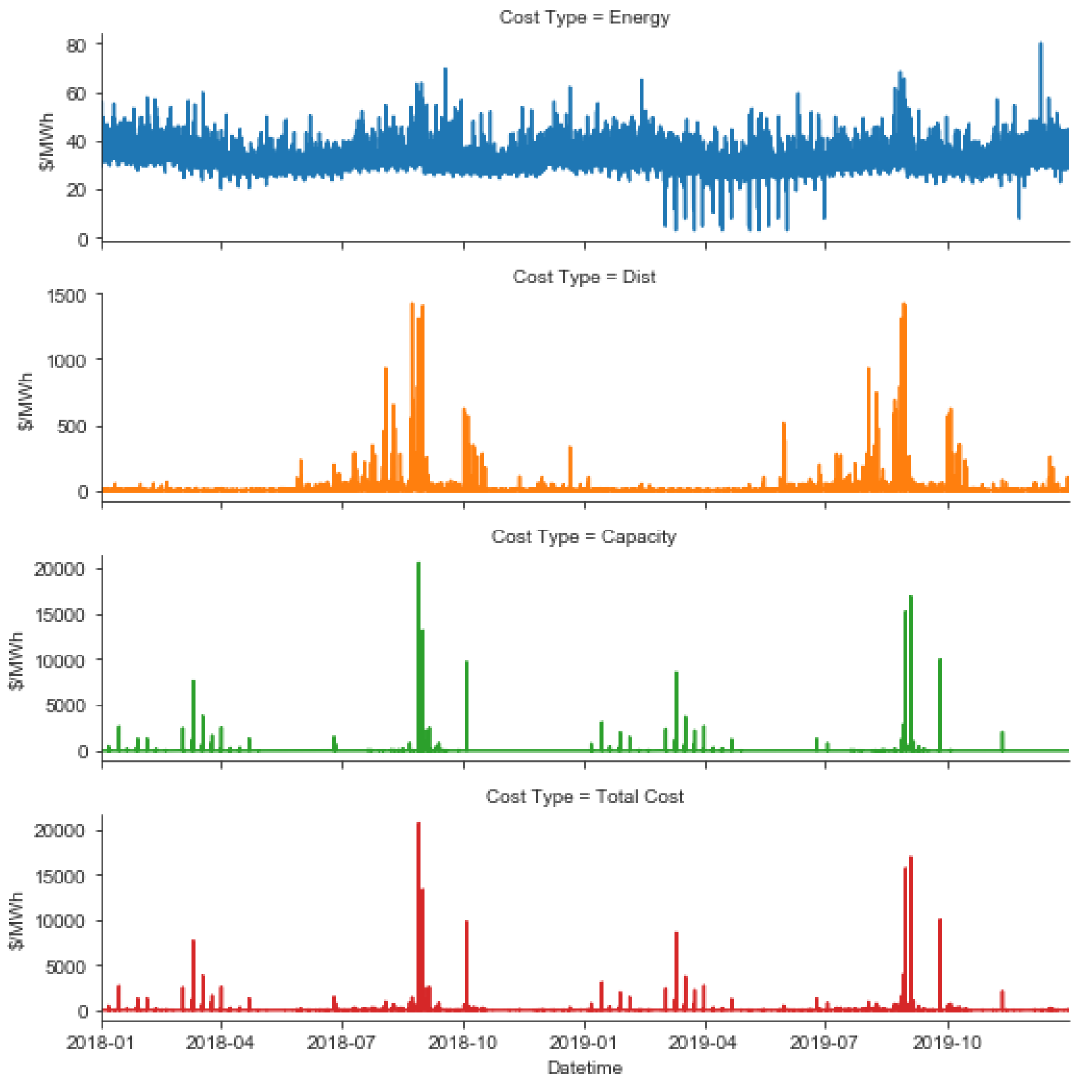

3.2 The Optimal 2-Season / 2-Period TOU Rate

Starting with a 2-season / 2-period TOU specification, I initialized the search algorithm 1,000 times. I also added two constraints: first, that each period be at least 4 hours in length and, second, that each season be at least three months long. Each initialization had a different, randomly selected starting seasons and starting periods. Across these 1,000 trials, the algorithm:

- Made 24,723 discrete steps and nearly 150,000 unique fit assessments.
- Took an average of 24.7 steps for a trial to reach its local minimum.
- Found the optimal solution in 19% of trials, suggesting both that the optimal solution is stable, but also that that TOU surface is riddled with non-convexities.[5] Adjacent solutions

[5] An investigation into the suboptimal trials suggests that many are orthogonally adjacent, or near adjacent, solutions along a ridge. Ridges are a challenging problem for hill climber algorithms, because hill climbers only adjust one element in the vector at a time, each step will move in an axis-aligned direction. If the target

shifting summer peak period forward one hour or the summer season forward one month are the next-best fitting TOUs and account for another 42% of solutions.

The optimal 2-season / 2-period is presented in Table 1.

Table 1: Optimal 2-Season / 2-Period TOU Specification & Prices

Season	Period	Hours (HE)	Percent of Year	TOU Price ($/MWh)	TOU Differentials
Summer	**Peak**	17-20	4.2%	$427.38	7.65 : 1
(Aug-Oct)	**Off-Peak**	All Others	21.0%	$55.89	
Winter	**Peak**	18-21	12.5%	$91.08	2.64 : 1
(Nov-Jul)	**Off-Peak**	All Others	63.3%	$34.53	

These price differentials are material: 7.65:1 in the summer and 2.64:1 in the winter and would likely drive a change in consumption patterns. While the peak prices are high, they are incurred for only a short duration. The summer peak period accounts for 4.2% of all hours and the winter peak period just 12.5%. More than 80% of hours have the low off-peak prices. (Note, extending the minimum duration of the peak-period would push these differentials down, if regulators deem these price differentials too extreme for policy reasons. Reduced price differentials would come at the expense of efficiency.)

Figure 2 depicts how the TOU rates compare to the underlying real-time prices. The figure sorts the dataset first by the TOU rate (highest to lowest) then by the real-time price. The result is a series of price-duration curves, conditioned on TOU price. Note that for legibility reasons, the results are plotted on a log scale.

function creates a narrow ridge that ascends in a non-axis-aligned direction, then a superior orthogonal move requires two inferior moves (one up, one over). Often, the hill climber will stall-out in these near-adjacent locations.

Figure 2: Distribution of Hourly and TOU Prices, sorted ($/MWh)

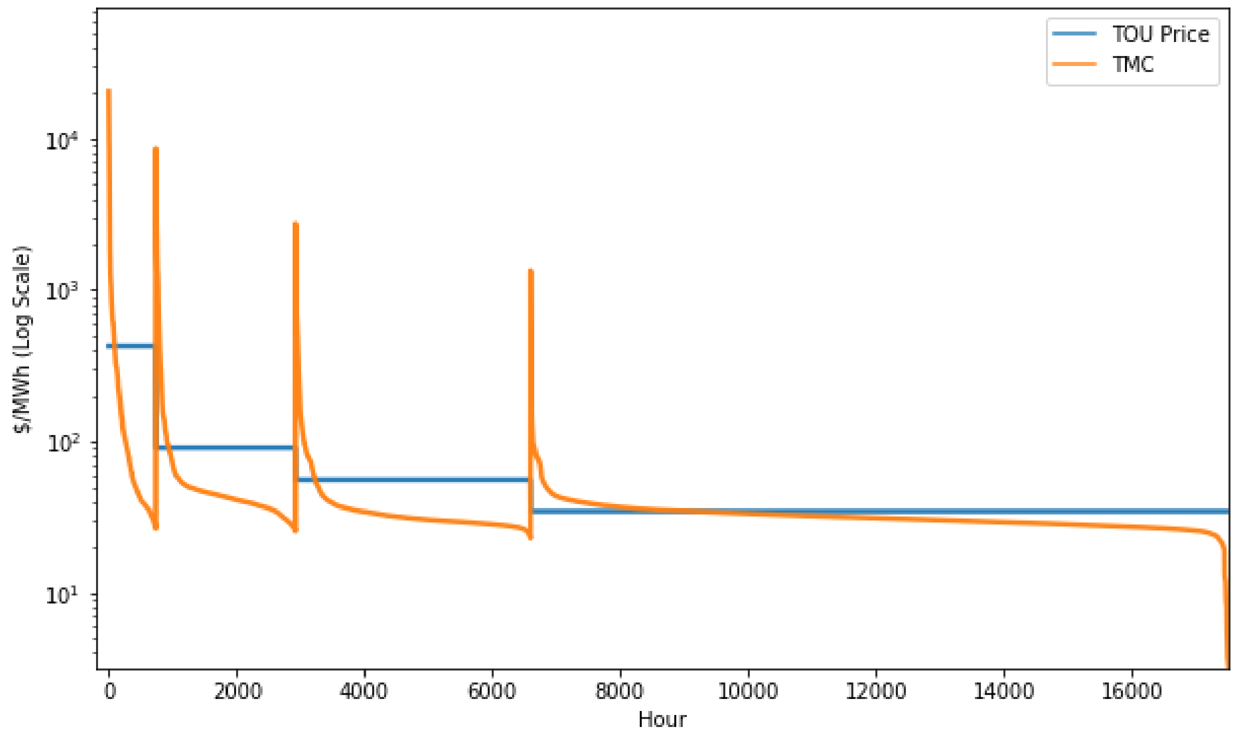

The all-in rates in Table 2 can be decomposed into various rate-components by calculating the load-weighted price of each component by season and period. For example, the overall costs of the 2-Season / 2-Period TOU rate can be split into discrete energy, capacity, and T&D cost components. Table 3 depicts these individual cost components.

Table 2: Cost Components of 2-Season / 2-Period TOU Tariff

Season	Period	TOU Price ($/MWh)	Energy ($/MWh)	Capacity ($/MWh)	Distribution ($/MWh)
Summer	**Peak**	$427.38	$37.40	$278.05	$111.92
(Aug-Oct)	**Off-Peak**	$55.89	$31.07	$6.17	$18.64
Winter	**Peak**	$91.08	$37.55	$40.81	$12.72
(Nov-Jul)	**Off-Peak**	$34.53	$30.39	$0.64	$3.50

This decomposition allows for regulators to appropriately group costs together into the standard bill components, rather than presenting a single, top-line rate. When considering these sub-components, it becomes apparent which subcomponents drive the all-in costs. Capacity and distribution costs drive costs in the summer peak period, accounting for 91% of the overall peak-period price. These capacity and distribution costs are 85% lower in the winter season, resulting in a much lower overall winter peak period rate. In the off-peak period of nine winter months, the $34.53/MWh all in rate is predominately the cost of energy. Distribution costs are non-zero in the winter off-peak period, but very small (about 10% of the winter off-peak rate).

3.3 Would a different Specification be Better? The Optimal 2-Season / 3-Period TOU Rate

Would a 2-season / 3-period specification be preferable to the 2-season / 2-period developed in the previous section? Using the same underlying dataset, I develop best-fitting 2/3 alternative rate then compare it to the 2/2 using the same method described in Section 2.5. Table 3 presents the tariff's optimal all-in rates by season and period. Note that the peak-periods are the same as in the 2/2 rate, but that the off-peak rate has been split into two: a new middle price adjacent to the peak period with a differential of about 1.69:1 in the summer and 1.21:1 in the winter, and a very low off-peak price which is mostly entirely comprised of energy costs.

Table 3: Optimal 2-Season / 3-Period TOU Specifications & Prices

Season	Period	Hours (HE)	Percent of Year	TOU Price ($/MWh)	TOU Differentials
Summer	**Peak**	17-20	4.2%	$427.38	9.03 : 1
	Mid	13-16	4.2%	$79.77	1.69 : 1
(Aug-Oct.)	**Off-Peak**	All Others	16.8%	$47.33	
Winter	**Peak**	18-21	12.5%	$91.08	2.77 : 1
	Mid	14-17	12.5%	$40.00	1.21 : 1
(Nov-Jul)	**Off-Peak**	All Others	49.9%	$32.93	

Qualitatively, these period rates look materially different from one another (with the possible exception of the mid and off-peak winter prices). The short-duration peak periods are high and should encourage reductions in peak period consumption. (Based on prior econometric analysis, it seems doubtful that that winter-time mid-period will meaningfully change behavior compared to the off-peak period; both are low enough, however, to shift consumption away from the peak period.)

The rates remain relatively simple and the year-round peak period enhances the simplicity and understandability touted by Bonbright. Quantitatively, there is a modest improvement in the new 3-period rates. Table 4 computes the AIC and BIC metrics for both models, which allows us to directly compare the information density of the two TOU specifications. Recall at a lower AIC or BIC is better.

Table 4: AIC and BIC Metric by TOU Specification

Specification Metric	**2-Season / 3-Period** (a)	**2-Season / 2-Period** (b)	**Difference** (c) = (a) - (b)
AIC	205,171	205,172	-1
BIC	205,233	205,219	+14

In Table 4, we observe that both models are of roughly equal quality, and that the 3-Period model has slightly lower AIC and slightly higher BIC values than the 2-Period alternative. This suggests that improved descriptive power of three-period specification approximately offsets the incremental complexity. While it is not much worse than the two-period alternative, it is not obviously better, either.

3.4 The Optimal TOU Rate using Dynamic Specifications

As noted in Section 2.6, TOU specifications can be endogenized to search a whole range of possible TOU configurations – differing numbers of seasons and differing numbers of periods in each season, Here, instead of presupposing a specific TOU specification, I randomly initialize possible TOU specifications by using a set of simple rules:

- The TOU rate must have 1-3 Seasons;
- Each Season must be at least 3 months long;
- Each Season must have 1-3 Periods; and,
- Each Period must be at least 4 hours long.

This allows for TOUs ranging from a single year-round price up to a TOU with 16 discrete time/price segments. To assess a range of different TOU specifications initializations, and possible non-convexities within a given specification, I ran the search algorithm 2,000 times. As before, each initialization finds TOU with the lowest SSE value. To find the best TOU rate, the 2,000 are then ranked by AIC – to allow for direct comparison across different specifications. The most efficient rate has seven periods spread across three seasons and is presented in Table 5.

Table 5: Optimal TOU Specification & Prices

Season	Period	Hours (HE)	Percent of Year	TOU Price ($/MWh)	TOU Differentials
Season 1	**Peak**	18-21	5.6%	$73.57	2.12 : 1
(April – July)	**Off-Peak**	All Others	27.9%	$34.76	
Season 2	**Peak**	17-20	4.2%	$427.38	9.03 : 1
(August – October)	**Mid**	13-16	4.2%	$79.77	1.69 : 1
	Off-Peak	All Others	16.8%	$47.33	
Season 3	**Peak**	17-20	6.9%	$118.88	3.53 : 1
(November – April)	**Off-Peak**	All Others	34.5%	$33.70	

The prices in each season have meaningful differentials which could facilitate load shifting for flexible loads such as EV charging or space-heating. The AIC metric does a reasonable job of screening out solutions with immaterial price differentials. It is possible that further work could collapse each of the off-peak prices in Seasons 1 and 3 into a single value – further improving pricing clarity.

The best specification using the BIC metric is the same 2x2 TOU described in Section 3.2.

4. Conclusions

The preceding analysis offers a proof-of-concept for finding high quality TOU specifications using algorithmic techniques. I develop a method for calculating optimal TOU tariffs using granular price and load data, using a curve-fitting approach. These optimal TOU tariffs include explicit tradeoffs between different cost categories and between different season or period durations. This analysis also offers economic and statistical methods to choose between different specifications (e.g. 2-

period vs 3-period specifications). With the basic structure established, this work could be easily extended to capture added nuance. For example, it could be extended to multiple rate classes or to capture how consumers may respond to the price signals embedded in the TOUs (by incorporating the price elasticity of demand). The use of SSE and AIC/BIC metrics also provides a path forward to quantitatively evaluate the relative merits of alternative tariff proposals offered during a rate-case or other rate-design proceeding at a Public Utility Commission.

One insight from this TOU selection methodology and the underlying cost allocation is that in most hours of the year, the most efficient TOU rate should just slightly higher than the expected cost of energy. The optimal peak periods are short in duration and very high price. These short peak periods reflect the ground truth that many system costs are sized to meet peak demands. These time-differentiated rates would allow for selling of electricity in off-peak hours to a house, car, or heat-pump at a price much lower than average all-in prices. These low prices are unsubsidized and do not cause cost shift from other customers – not so long as the T&D and capacity cost allocation remains representative of system costs.

This work is of practical value because it provides an objective metric against which proposed rates can be evaluated. It also offers a path forward for modeling how retail rate designs should change in the face of increasing electrification of transportation and heating. For example, this work suggests that EV specific rate-classes could offer very low charging off-peak prices because these periods do not materially affect capacity or T&D costs. All-in rates measured below $60/MWh certainly look appealing in relation to the summer-peak rates nine-times that. It also hints that policymakers may need to seriously contemplate shaping T&D and capacity costs in the cost-allocation process as energy-only TOU rates will generally exhibit modest differentials. TOU rates offer a first-step towards more dynamic retail pricing schemes, and the methods outlined in this paper offer an objective method for deriving and defining the TOU rates.

---

[i] Severin Borenstein. 2005. "Time-Varying Retail Electricity Prices: Theory and Practice," in Griffin and Puller, eds., Electricity Deregulation: Choices and Challenges, Chicago: University of Chicago Press.
[ii] Severin Borenstein and Stephen Holland. 2005. "On the efficiency of competitive electricity markets with time-invariant retail prices." Rand Journal of Economics. Vol. 36, No. 3, Autumn 2005 pp. 469–493. Available at: http://faculty.haas.berkeley.edu/borenste/download/Rand05RTP.pdf
[iii] Andrew Satchwell, Peter Cappers, and Galen Barbose. 2019. "Current Developments in Retail Rate Design: Implications for Solar and Other Distributed Energy Resources". Energy Analysis and Environmental Impacts Division Lawrence Berkeley National Laboratory. Available at: https://emp.lbl.gov/publications/current-developments-retail-rate.
[iv] Peter Cramton, Axel Ockenfels, and Steven Stoft. 2013. "Capacity Market Fundamentals." Economics of Energy & Environmental Policy, Vol. 2, No. 2 at 1. Available at: http://www.cramton.umd.edu/papers2010-2014/cramton-ockenfels-stoft-capacity-market-fundamentals.pdf.
[v] Borenstein & Holland (2005).
[vi] Koichiro Ito. 2014. "Do Consumers Respond to Marginal or Average Price? Evidence from Nonlinear Electricity Pricing" American Economic Review. Vol. 104, No. 2 (FEBRUARY 2014), pp. 537-563 (27 pages).
[vii] Brendon Baatz. 2017. "Rate Design Matters: The Intersection of Residential Rate Design and Energy Efficiency". American Council for an Energy-Efficient Economy Report U1703. Available at: https://www.aceee.org/sites/default/files/publications/researchreports/u1703.pdf.
[viii] Thomas Chiles, Kenneth Shutika, and Philip Coleman (2015). "Riding the Electricity Market as an Energy Management Strategy: Savings from Real-Time Pricing". Available at: https://www.osti.gov/servlets/purl/1425679

[ix] Juan Pablo Carvallo, Nan Zhang, Sean Murphy, Benjamin Leibowciz, and Peter Larsen. 2020. "The economic value of a centralized approach to distributed resource investment and operation". Applied Energy Volume 269. Available at: https://doi.org/10.1016/j.apenergy.2020.115071.
[x] Kellen Schefter and Becky Knox. 2018. "Accelerating Electric Vehicle Adoption". Edison Electric Institute. Available at: https://www.eei.org/issuesandpolicy/electrictransportation/Documents/Accelerating_EV_Adoption_final_Feb2018.pdf, at 5-6.
[xi] Ryan Hledik, John Higham and Ahmad Faruqui. 2019. "Emerging Landscape of Residential Rates for EVs". Available at: https://www.fortnightly.com/fortnightly/2019/05/emerging-landscape-residential-rates-evs.
[xii] Paul Joskow and Catherine Wolfram. 2012. "Dynamic Pricing of Electricity". American Economic Review, 102(3): 381–385 http://dx.doi.org/10.1257/aer.102.3.381.
[xiii] Ryan Hledik, Ahmad Faruqui, and Cody Warner. 2017. "The National Landscape of Residential TOU Rates A preliminary summary". Available at: http://files.brattle.com/files/12658_the_national_landscape_of_residential_tou_rates_a_preliminary_summary.pdf, slide 5.
[xiv] Satchwell et al (2019), at Section 3.
[xv] Goutam Dutta and Krishnendranath Mitra. 2016. "A literature review on dynamic pricing of electricity." *J Oper Res Soc* **68,** 1131–1145. https://doi.org/10.1057/s41274-016-0149-4https://link.springer.com/article/10.1057/s41274-016-0149-4
[xvi] Ahmad Faruqui. 2018. "Estimating the impact of innovative rate designs. Available at: http://files.brattle.com/files/13955_estimating_the_impact_of_innovative_rate_designs.pdf.
[xvii] Jonathan Cook, Candice Churchwell, and Stephen George. 2014. "Final Evaluation for San Diego Gas & Electric's Plug-in Electric Vehicle TOU Pricing and Technology Study". Nexant, Inc. Available at: https://www.sdge.com/sites/default/files/SDGE%20EV%20%20Pricing%20%26%20Tech%20Study.pdf, Section 5.
[xviii] James Bonbright. 1961. ""Principles of Public Utility Rates" Columbia University Press. Available at: http://media.terry.uga.edu/documents/exec_ed/bonbright/principles_of_public_utility_rates.pdf, at Chapter 16.
[xix] Jim Lazar and Wilson Gonzalez. 2015. "Smart Rate Design for a Smart Future." Regulatory Assistance Project. Available at: https://www.raponline.org/knowledge-center/smart-rate-design-for-a-smart-future/.
[xx] Carl Linvill. 2018. "Rate Design to Maximize Grid Benefits: Smart EV Rate Design is Smart Rate Design". Regulatory Assistance Project. Presented at CPUC ZEV Rate Design Forum. Available at: https://www.raponline.org/wp-content/uploads/2018/06/rap_linvill_cpuc_zev_rate_design_2018_june_7.pdf
[xxi] James Sherwood et al. (2016). "A Review of Alternative Rate Designs: Industry experience with time-based and demand charge rates for mass-market customers. Rocky Mountain Institute. Available at: http://www.rmi.org/alternative_rate_designs.
[xxii] Melissa Whited, Avi Allison, and Rachel Wilson. 2018. "Driving Transportation Electrification Forward in New York
Considerations for Effective Transportation Electrification Rate Design". Prepared for Natural Resources Defense Council. Available at: https://www.synapse-energy.com/project/best-practices-electric-vehicle-rate-design.
[xxiii] For example, *see* California Public Utility Commission Decision 15-07-001 (July 3, 2015), section 6 on TOU rates. Available at: https://docs.cpuc.ca.gov/PublishedDocs/Published/G000/M153/K110/153110321.PDF. "It has been said that rate design is both a science and an art. For a default TOU rate to be successful, the design should be based on empirical evidence that supports both measurable benefits of TOU on the grid, and the acceptance and understanding of TOU rates by the residential customer." Decision at 130.
[xxiv] Russell, S. J., & Norvig, P. (2004). Artificial Intelligence: A Modern Approach. Upper Saddle River, NJ: Prentice Hall, at Chapter 4.
[xxv] Lazar and Gonzalez (2015), at 44.
[xxvi] Sherwood et al. (2016), at 25.
[xxvii] Bonbright (1961) at 291: "The 'practical' attributes of simplicity, understandability, public acceptability, and feasibility of application".
[xxviii] This violates another of Bonbright's criteria (p292ff).

[xxix] Cost allocation decisions fall well outside the scope of the present analysis. *See* Lazar, J., Chernick, P., Marcus, W., and LeBel, M. (Ed.). (2020, January). Electric cost allocation for a new era: A manual. Montpelier, VT: Regulatory Assistance Project. Available at: https://www.raponline.org/wp-content/uploads/2020/01/rap-lazar-chernick-marcus-lebel-electric-cost-allocation-new-era-2020-january.pdf

[xxx] Southern California Edison (SCE) GRC Proceeding (Phase I). CPUC Docket # A.16-09-003.
Load Data: https://www.sce.com/regulatory/load-profiles/dynamic-load-profiles
Price Data: Obtained through conversation with SCE.