The wholesale market value of dispatchable efficiency for commercial air conditioning

Jonathan Woolley [1,2], Thomas Jawin [3]

1. *University of California, Berkeley. Center for the Built Environment.*
2. *University of California, Davis. Western Cooling Efficiency Center*
3. *University of California, Berkeley, Civil & Environmental Engineering*

Semester Project - ENE 254 Electric Power Systems, December 15, 2015

ABSTRACT

This project develops the business value proposition for using a dispatchable efficiency measure for commercial air conditioning as a type of demand response. We explore this opportunity for an efficient air conditioning technology that reduces peak electricity demand from cooling by 40%. We argue that:

- Dispatchable efficiency is distinct from other types of demand response
- Dispatchable efficiency should be valued in wholesale markets
- Demand side reductions should be assigned as the preferred resource
- Operators of dispatchable efficiency should be compensated according to the marginal benefit they provide for other consumers (we assume this is equivalent to the market clearing price)
- To the extent that a portfolio of dispatchable efficiency is reliable and actively controllable, an operator should also be compensated for ancillary services such as spinning reserve

We construct a hypothetical business model and market operating scheme involving commercial building owners, a load serving entity, an efficiency operator (aggregator) and the wholesale electricity market. We develop an hourly annual model for the demand reduction that would be generated by a regional portfolio of the dispatchable air conditioning efficiency measure, and we use an annual history of market clearing prices from CA-ISO to estimate the revenue that could be obtained from selling demand reduction on the wholesale market. The model is also used to estimate the value of reduced retail energy purchases for the end user and reduced monthly demand charges.

The annual revenue estimate is used to consider a business structure in which:

1. The efficiency aggregator pays all capital expenses for installation of the efficiency measure
2. The building owner pays zero (or reduced) capital for the upgrades
3. The utility charges the building owner according to normal rate schedules
4. The building owner benefits from reduced energy expenses ($/kWh) and demand charges ($/kW)
5. The building owner's utility bill savings are shared with the aggregator
6. The efficiency aggregator pays all operating costs for the efficiency measure
7. The aggregator uses web-connected controls to manage and control the measure
8. The hourly demand reduction is aggregated across a regional portfolio of buildings.
9. The aggregator is compensated at the wholesale market clearing for the demand reduction generated each hour

Furthermore, we compare this business structure to similar variations that assign costs and benefits in different ways. We also compare it to the standard equipment sale-and-services approach that is currently used for the technology. We show that the strategy can generate an annual ROI > 100%, depending on how capital expenses and the electricity bill savings are divided between the building owner and the efficiency aggregator. Lastly, we explore the strategic advantages for each of these business structure variations compared to other types of energy services agreements and efficiency financing strategies.

INTRODUCTION & PROBLEM EXPLORATION

Demand response reduces load by temporarily cutting unused energy services

Traditionally, demand response has been characterized by controls that shut off non-critical services for periods that are deemed strategic by the utility. In buildings, a demand response event could trigger lights to switch off or dim, and may turn off air conditioning or reduce cooling capacity for a period of time. These changes are executed in ways that do not interrupt critical operations, but inevitably the demand reduction is achieved by reducing the level of service provided. Advocates for these strategies argue that the reduced service is acceptable to end users. However, if reduced lighting density and warmer temperatures are acceptable, these strategies ought to be adopted as permanent energy reducing measures. After all, the possibility that energy services can be reduced without notice indicates that the energy use is not providing utilized benefits.

The value of efficiency is not fully recognized by wholesale markets

According to FERC Order 745, the demand reduction from temporarily eliminating load can be compensated at the current locational marginal price, through day-ahead and hour-ahead wholesale electricity markets. Unfortunately, if an energy-conscious building owner were to permanently eliminate the some under-utilized energy consumption, they would not be compensated in the same way. For example, if a building owner were to permanently de-lamp over-lit corridor spaces, their action would have benefits to other energy end users - because the reduced demand would shift the wholesale marginal cost of electricity - but the building owner would not be compensated for these benefits. As long as the marginal price of electricity increases monotonically with respect to load, the de-lamping project would have benefit to other customers at any time. Moreover, it would have a larger benefit during higher load periods when the marginal price of electricity is more sensitive to load. Despite this fact, permanent load reduction is not given value on the wholesale market in the same way as demand response.

There are also many efficiency improvements for buildings that reduce energy consumption without reducing the level of service. The demand reduction from efficiency benefits all energy end users because it permanently shifts the wholesale marginal cost of electricity. Efficiency can have greater impacts overall because, contrary to demand response, it is not limited to affecting non-critical under-utilized services. Instead, efficiency measures can provide equal service with less energy use. Nevertheless, efficiency is not recognized in wholesale electricity markets.

The effects of energy efficiency cannot be measured reliably

There are good reasons that efficiency and permanent load reduction are not valued in wholesale electricity markets. Most importantly, the energy savings associated with these activities cannot be measured because they only represent the difference from a hypothetical energy use. Energy savings can only be estimated by comparing actual energy use to an expected baseline energy use. However, in most scenarios the energy consumption at any point in time is influenced by many uncontrolled variables, giving baseline projections a high degree of uncertainty. Moreover, future changes such as tenant improvements, activity reprogramming, or additional efficiency upgrades have complex interacting effects that make it incredibly difficult to confidently attribute particular energy savings to the original efficiency measure. As efficiency standards and construction practices change, it can be argued that previously installed efficiency measures will no longer have impact because the hypothetical baseline scenario would have evolved naturally over time.

Despite the baseline uncertainty, there are many instances where efficiency and permanent load reduction have demonstrated market value. The energy and demand savings are at least reliable enough to motivate some building owners to invest in efficiency on their own. Energy services companies find that baseline energy use can be estimated reliably enough to guarantee energy savings by contract with their customers;

although they usually also establish agreements that allow baseline projections to be re-assessed over time. In some regions, demand side efficiency has even been valued in forward capacity markets. For example, ISO-NE has allowed demand management products to qualify as capacity resources (CA ISO 2011). The nation's first Energy Efficiency Utility, Efficiency Vermont, operates largely on selling demand reduction from end use efficiency in this market (Efficiency Vermont 2015). However, the value assigned in this forward capacity market is mostly associated with deferred public investments in transmission and generation, not with the benefits surrounding procurement of electricity in spot markets at wholesale locational marginal prices. It is clear that end use efficiency measures and permanent demand reduction do have benefits, but they are currently not recognized in wholesale electricity markets.

Controls could allow certain efficiency measures to participate reliably in spot markets

Demand response is uniquely recognized in day-ahead and hour-ahead electricity markets because it can be dispatched and controlled in real time. The primary advantage is that the impacts of demand response can be observed immediately and assigned a reliable value - the demand reduction can be measured with acceptable confidence. However, the same capabilities could be achieved by "dispatchable efficiency", where certain technologies could be controlled in real time to establish estimates of demand reduction with a degree of certainty that is equal to simple load shedding strategies. Dispatchable efficiency is only appropriate for a certain class of efficiency measure that can be controlled to revert to a baseline mode. Improved control sequences like occupancy responsive lighting, demand controlled ventilation, or temperature and pressure reset improvements could all be deployed as dispatchable efficiency. Physical improvements such as building envelope insulation and fenestration upgrades would not be suitable.

Measures deployed as dispatchable efficiency could operate to deliver energy savings at all times, but occasionally they would be disabled temporarily to establish a new baseline. As a result, the approach would be resilient to external changes such as occupant behavior, or physical improvements in a building, or upgrades for other energy end uses. Customers would benefit more from dispatchable efficiency than they do from typical demand response because their energy use would be reduced at all times and their monthly demand charges would be lower. The customer benefits of demand response have so far been restricted to limited events when the utility deems that demand response is strategic. Arbitrated by the utility alone, these events are largely determined to hedge against un-forecasted price spikes. Since retail rates are regulated by public utilities commissions through multi-year rate case scenarios, it is arguable whether or not the benefits of these demand response events are fully realized by electricity consumers. Moreover, dispatchable efficiency benefits all energy end users because it permanently shifts the wholesale marginal cost of electricity. Conventional demand response has a more limited benefit because it is only triggered during the most significant price spikes.

Exploring the implications and business value proposition for dispatchable efficiency

FERC Order 745 has laid the foundation for demand side resources to participate directly in wholesale spot electricity markets. The order directs that for any hour where the total benefit of demand response is larger than the cost to procure those resources, demand reduction should be compensated at the current locational marginal price for electricity generation. The order is especially significant because it allows third party demand resource aggregators to participate directly in wholesale electricity markets. This bypasses utilities and the conflicting incentives associated with their regulated business model.

In this paper, we explore the dispatchable energy savings that could be achieved from a commercial air conditioning efficiency measure referred to as dual-evaporative pre-cooling. The retrofit is unique in that it can be added as a wrapper to an existing rooftop air conditioner without interfering with the internal controls and operation of that machine. When the retrofit is enabled it improves the operating efficiency of the existing system, and when it is disabled the equipment effectively reverts to baseline performance. This commercially available measure incorporates web-based controls that allow for remote monitoring, fault detection, and dispatchable operation. We explore the business opportunities for the technology where demand reduction in every hour would be aggregated across a portfolio of commercial buildings and sold as a demand side resource on the wholesale electricity market.

We developed a simplified model to estimate the energy savings for dual-evaporative pre cooling in any location for every hour in the year. The model uses historical climatological data, and the coincident locational marginal price for electricity, to assess the value of demand reduction on the wholesale spot market. It also allows for input of retail electricity rate schedules and demand charges in order to calculate the financial benefits for end users who employ the efficiency strategy. Furthermore, we obtained information about capital costs and operating expenditures from the technology manufacturer in order to develop a complete financial assessment that considers debt service on capital expenses, as well as the costs of service, warranty, maintenance, and web-based controls operation.

Results from the model are used to evaluate the value propositions for five different business structures including the business model currently employed by the technology manufacturer. Aside from the fact that the current business model does not derive value from the wholesale spot market, the models explored differ mainly in the way that they assign costs and benefits to a customer and a demand resource aggregator. Beyond the quantitative financial assessment, we also discuss some of the strategic advantages and disadvantages for each business opportunity, compared to other types of energy services agreements and efficiency financing strategies. For example, in many of the approaches the wholesale value of demand reduction could enable the measure to be installed for zero first cost to the customer. This would bypass challenges with access to capital, circumvent split incentives problems, and ease issues with timing mismatch – all phenomena that stifle the economically rational uptake of efficiency measures.

TECHNOLOGY OVERVIEW

Dual-evaporative pre-cooling as an example of dispatchable efficiency:

Dual-evaporative pre-cooling is one design approach for hybrid air conditioning. Hybrid air conditioners, in general, use strategic combinations of components to cool and heat buildings more efficiently than vapor compression alone. Hybrid air conditioners might incorporate desiccant dehumidification, heat recovery ventilation, indirect evaporative cooling, or other modes. Most commercially available hybrid systems also utilize vapor compression, but use other more efficient methods to accomplish a significant portion of the building cooling requirements, or to improve the efficiency of the vapor compression cycle. Dual evaporative pre-cooling is a unique technology that allows for retrofit transformation of existing air conditioners into hybrid systems. The strategy uses direct evaporative cooling to pre-cool condenser air. Then, cooled water from this process is collected and circulated through a water-to-air heat exchanger that cools ventilation air. The water absorbs heat from the ventilation air and is recirculated to and cooled through the evaporative process at the condenser inlet. This coupled process creates a preliminary stage of cooling, while simultaneously decreasing compressor pressure and the amount of work necessary.

The dual-evaporative pre-cooler can be deployed as a retrofit for any existing air conditioner without disturbing the internal controls for the existing system, and without requiring integration with the whole building sequence of operations. Since the technology is fully self-contained, it does not require substantial adaptation for different buildings with different controls infrastructure. The system incorporates a web-based monitoring system that is currently used for fault detection. Through this project we envision that the current monitoring system would also be used to control dual-evaporative pre-cooling as a dispatchable efficiency measure. The technology is uniquely suited to this strategy because it is installed as a wrapper for existing air conditioners - when it is disabled the existing system reverts to baseline operation. Although we do not consider it in our model, the dispatch strategy for a regional portfolio of pre-coolers could be designed to operate at partial savings for certain periods in order to also provide ancillary services equivalent to spinning reserves, or frequency control.

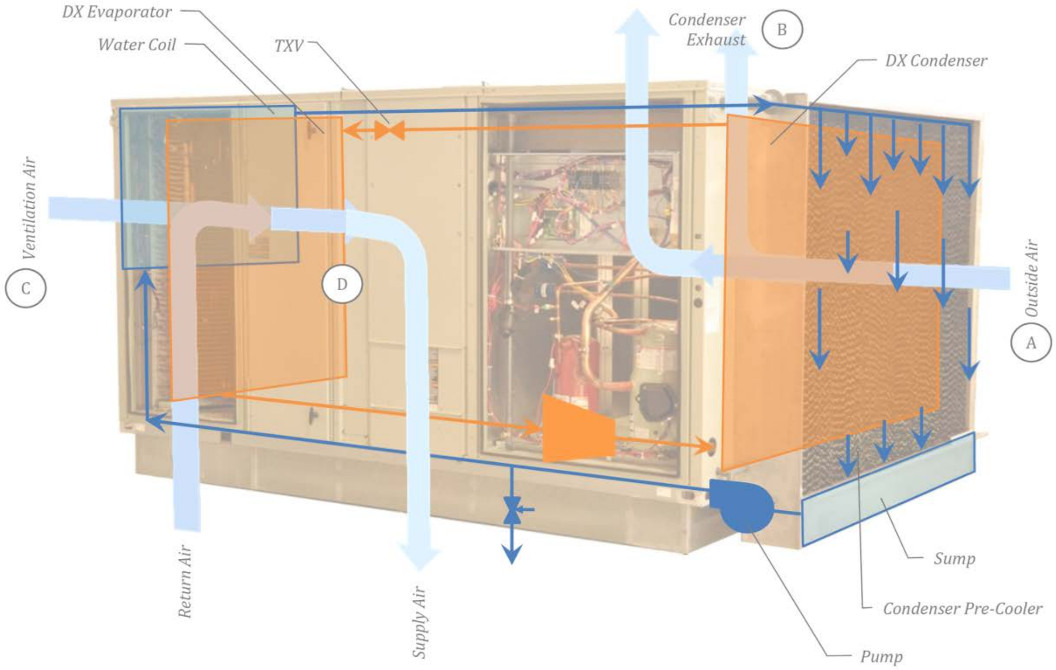

A. Hot dry outside air is drawn through a fluted cellulose evaporative cooler located at the inlet of the vapor compression condenser coil. Water is delivered through a manifold at the top of the media, and flows through the fluted channels in contact with airflow. The air and water are both cooled by evaporation, and excess water drains by gravity to a stainless steel sump.
B. Cool moist air is drawn across the condenser coils then exhausted from the equipment through condenser fans. When operating in a vapor compression mode, heat is rejected to this airstream, but the fans can also operate independent of compressors to cool water for ventilation air cooling.
C. Water that drains from the evaporative cooler is collected in a sump, then circulated through a water coil located at the ventilation air inlet to cool outside air before it enters the vapor compression evaporator.
D. Return air and outside air are cooled at the evaporator coil. Capacity and cooling efficiency for the vapor compression circuit both increase as a result of the condenser pre-cooling and ventilation pre-cooling.

Figure 1: Conceptual schematic for rooftop unit with dual-evaporative pre-cooling

Previous laboratory and field studies of this hybrid air conditioning strategy have concluded that it can reduce energy use for cooling by more than 40% at peak compared to conventional minimum efficiency standard commercial air conditioners. Figure 2 plots the coefficient of performance observed for one unit studied in Ontario, California (Woolley 2015). The performance of this unit serves as the basis for our analytical model. For consistency, the annual scenario simulated for this report uses historical climatological data and day ahead wholesale locational marginal price data for the Ontario region.

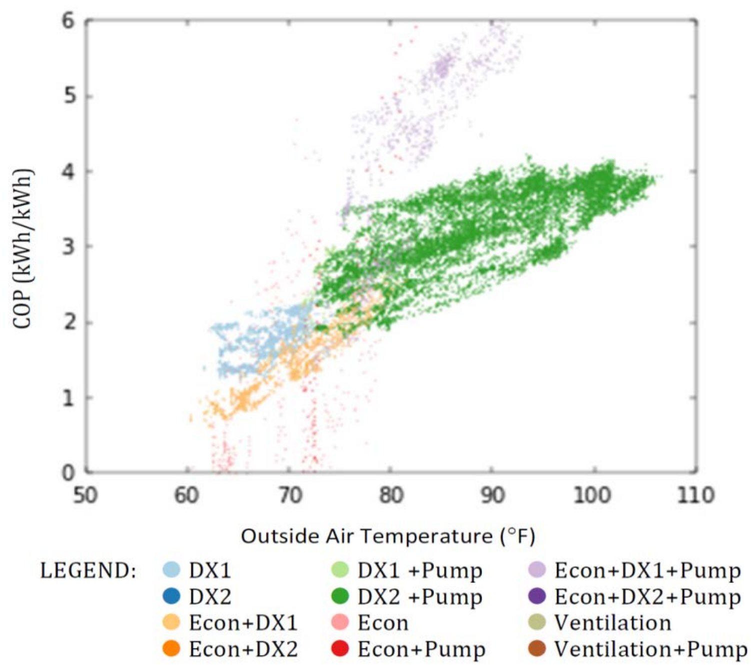

Figure 2: Sensible System Energy Efficiency Ratio for a rooftop unit with dual-evaporative pre-cooling in each mode of operation (Woolley 2015)

METHODS:

Explanation of the analytical model:

We developed a spreadsheet model to estimate energy savings for dual-evaporative pre cooling in any location for every hour in the year. The model is structured in a generalized way so that a user can input the following information in order to simulate any scenario that is desired:

- Annual hourly outside air temperature data
- Annual hourly locational marginal price for electricity
- Hourly electricity rate schedule and monthly demand charges
- Building operating schedule
- Building sensible thermal load characteristics
- Baseline equipment efficiency characteristics
- Equipment efficiency characteristics with dual-evaporative pre-cooling
- Capital costs for dual-evaporative pre-cooling
- Annual service and operating costs for dual-evaporative pre-cooling
- Financing terms

Further, the model summarizes results to develop an annualized value proposition for several different building models. The following sections discuss the assumptions underpinning our model, the specific structure of each major component in the model, and the sample inputs that were used to develop the simulation results presented in this report.

Model assumptions:

We constructed hypothetical scenarios where the demand reduction from dual evaporative pre-cooling is sold in wholesale electricity markets as a type of demand-side resource. Our models assume that dispatchable efficiency is compensated for demand reduction on an hourly basis, for every hour in the year where it generates savings, at parity with the concurrent wholesale market value for generation.

According to FERC Order 745, demand reduction should only be dispatched during those periods when compensation at the current market clearing price generates net benefits to all users. Our approach conflicts with FERC 745 since we do not perform a "net benefits test". (Xu 2011) However, the most significant savings from dual-evaporative pre-cooling are coincident with high demand periods when net benefits would be higher. Furthermore, by the same logic that underpins FERC's net-benefits test we would argue:

- Demand reduction should be compensated at all times, but only at a rate that is equivalent to the benefits that it produces
- When the rate of increase for marginal cost is low, demand reduction should be valued lower than the marginal cost of electricity
- When the rate of increase for marginal cost is high, demand reduction should be valued higher than the marginal cost of electricity

Instead of instituting a more dynamic logic that recognizes these ideas - where the rate of compensation would depend on the current market clearing price, the rate of change in price as load increases, and the total load on the grid - we used the historical LMP data as the value for all hours where there are savings.

Furthermore, we also assume that customer demand charges would persist in their current form, despite the fact that demand reduction is also valued on the wholesale market. This assumption is based largely on the idea that demand charges are designed to reflect the degree to which customers disproportionately contribute to the need for overall transmission and distribution infrastructure. However, even if demand charges are intended to give economic signals that are associated with wholesale market prices, we argue that demand reduction should still result in both reduced charges and compensation from the market for the capacity.

Figure 3 (on the next page) demonstrates why this is the case at when demand and marginal cost are at the threshold of net-benefits.

Between case-A and case-B marginal price increases in response to an increase in load. At this point, customers should be indifferent about whether or not the price they pay supports increased generation (as in case-B) or instead supports some customers within their cohort to reduce demand (as in case-C). The customers who reduce demand are compensated at the market clearing price: the same rate at which generators are compensated for positive generation. Also, these customers do not have to pay for the energy they do not use, and would not have to pay demand charges for power they do not require. Meanwhile, other customers pay the same amount they would pay without demand response, and generators are compensated at a fair rate for the energy they generate.

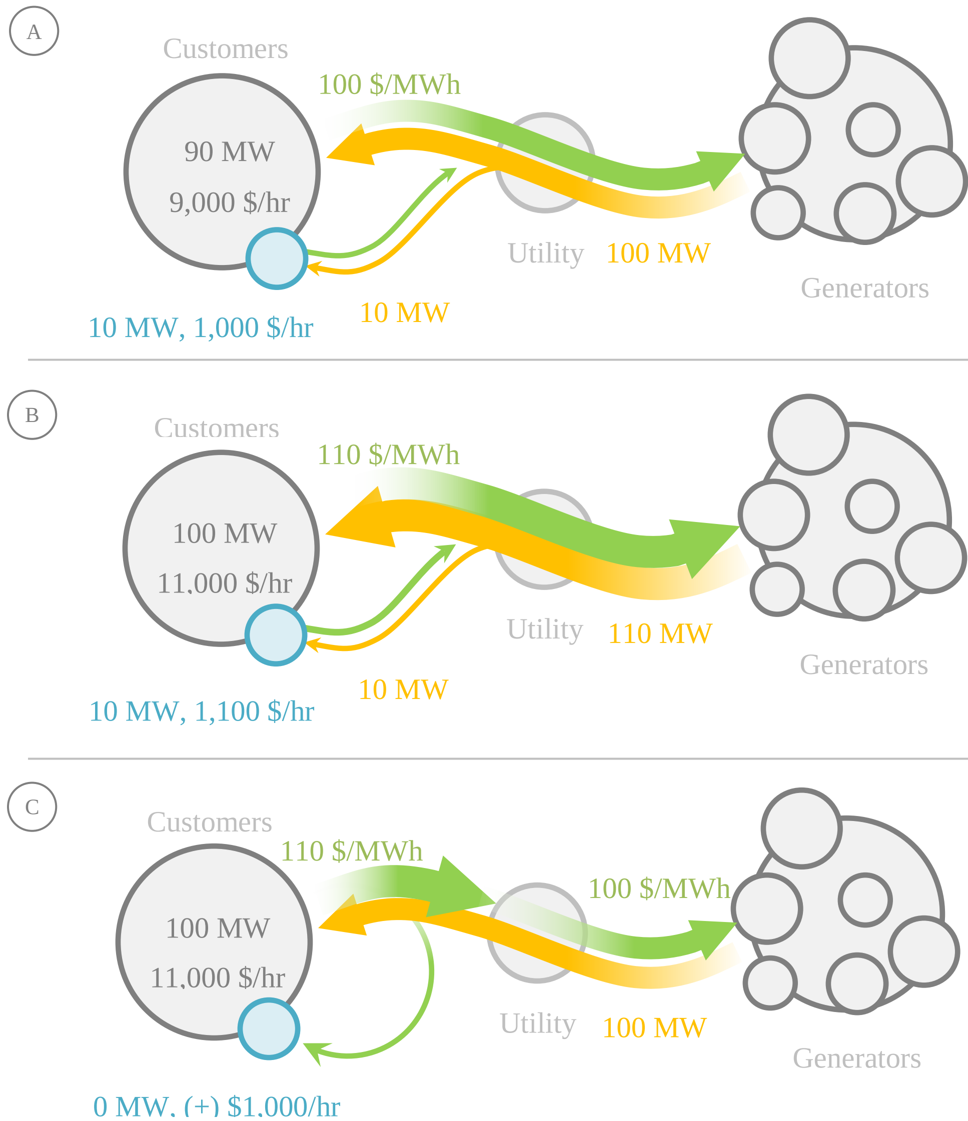

Figure 3: The tradeoff for demand response at the net-benefits threshold price

Model for thermal load:

The model is constructed to assess the value of demand reduction for dual-evaporative pre-cooling applied to 1000 nominal tons of conventional air conditioning equipment. This is a manageable scale to consider in the model, and it might not be worthwhile to aggregate savings from a smaller portfolio. The scale is also achievable and would produce considerable annual benefits. 1000 tons represents the air conditioning load for about 10 large format retail stores, or between 20-50 small commercial businesses.

The floor area served by each ton of air conditioning changes by application and by climate, and the sensible thermal load profile in each building is also different. In order to accommodate the range of possibilities, our model is constructed to allow input of any load profile. For the simulation results presented in this report, we developed a hypothetical piecewise-linear single-breakpoint mode for cooling load (kBtu/hr) as a function of outside air temperature. We developed two thermal load profiles: one for

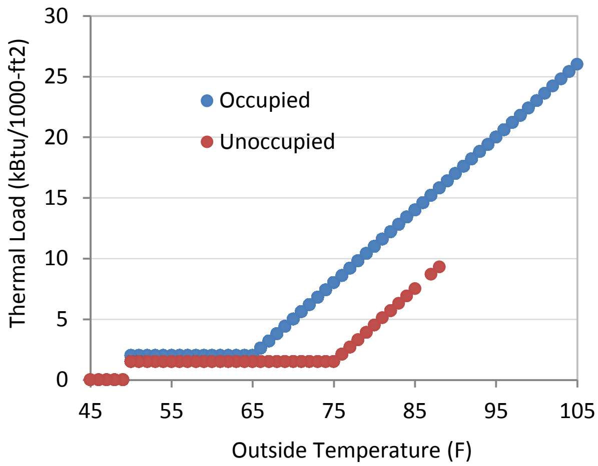

Figure 4: Thermal load as a function of outside temperature for occupied and unoccupied periods

occupied periods and one for unoccupied periods. We assumed that the occupied sensible thermal load at 95°F outside air temperature is 20 kBtu/1000 ft^2, which equates to roughly 600 ft^2/ton, and 120 cfm ventilation air per ton in a building with 0.2 cfm ventilation per ft^2. Below 65°F, the occupied thermal load was fixed at 2 kBtu/ft2, and the unoccupied thermal load was fixed at 1.5 kBtu/ft2. Below 55°F the thermal load was reduced to zero. However, any condition below 70°F has no bearing on the model results since dual-evaporative pre cooling does not have an effect below that point. Figure A presents the simulated thermal load in occupied and unoccupied periods. Figure 5 (A) plots a time series trend for thermal load for one sample week in our simulation, and Figure 5 (B) plots the sum of thermal load for each month in the annual simulation.

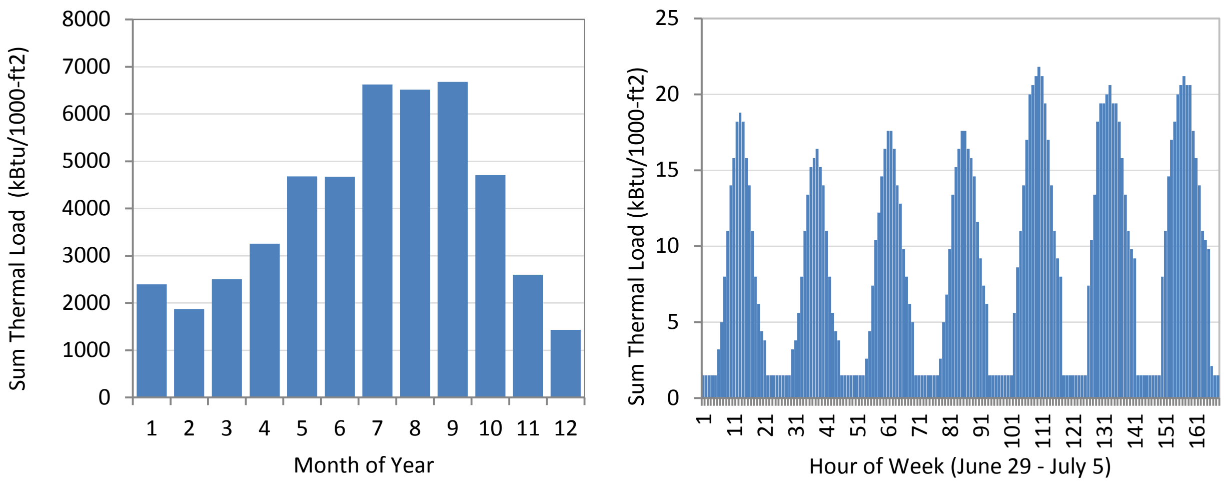

Figure 5 (A:) Time series of thermal load for one week, Figure 5(B): Sum of thermal load in each month for one year in Ontario CA

Model for cooling efficiency to determine electric energy use and savings:

The sensible thermal load in each hour was divided by models for cooling efficiency to determine the electrical energy use in each hour. Our models of sensible efficiency for conventional air conditioning, and for hybrid air conditioners with dual-evaporative pre cooling are plotted in in Figure 6 (A). Both models are defined as linear functions of outside air temperature; the model for performance with dual-evaporative pre-cooling uses a piecewise linear model with a breakpoint at 70°F, where Dual-Evaporative Pre-Cooling begins to improve efficiency. The two models have identical efficiency up to 70°F, but the model for units with dual-evaporative pre-cooling predicts significantly higher efficiency as temperatures increase. We presented an explanation for these efficiency advantages in section "Technology Overview".

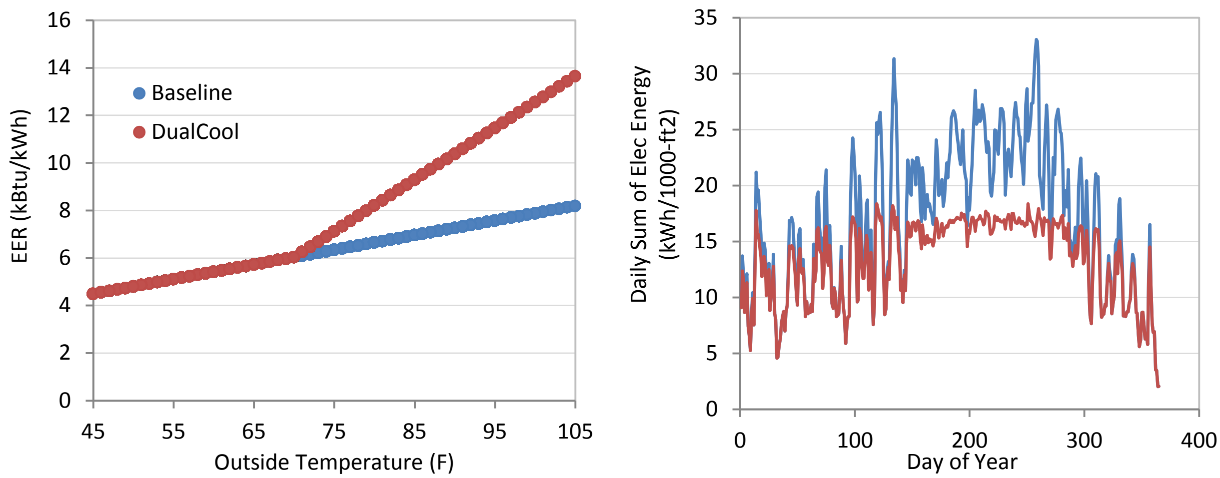

Figure 6: (A) Sensible System Energy Efficiency Ratio for a rooftop unit with and without dual-evaporative pre-cooling (B) Daily sum of electric energy consumption for cooling, with and without dual evaporative pre-cooling, for one complete year in Ontario California

The resulting electric energy consumption in in each scenario is plotted in Figure 6 (B) as a time series trend for every hour of the year. This plot illustrates that the measure has very little impact during milder months, but that it provides substantial energy savings during warmer periods. As a result, the total annual savings is largely dependent on climate.

The hourly energy savings data illustrated in Figure 6 (B) is multiplied by an hourly electric rate schedule to determine hourly savings in energy costs, and the same data is multiplied by the concurrent locational marginal price in order to determine the wholesale market value. The difference in each month's maximum hourly demand is used to assess reduction in demand charges - except that we assume that the maximum 15-minute-average demand is 25% larger than the month's maximum hour-average demand. For the results presented in this report, we assumed a retail electricity rate of 0.12 $/kWh - which is low for many large commercial customers in California. We assessed a monthly demand charge of $10/kW.

Wholesale value for demand reduction:

Wholesale value for electricity is determined as a locational marginal price, which reflects both economic and physical characteristics of the electrical network. According to the logic established by FERC 745, demand side resources should be compensated for reducing load at the current locational marginal price. Since the locational marginal price varies for different nodes throughout a region, demand reduction in some areas would be more cost effective than in others. In this project, we assume that savings from dual-evaporative pre-cooling is compensated in any hour that the efficiency measure provides savings. We downloaded data from CA ISO's Open Access Same-Time Information System (OASIS) for day ahead locational marginal price on a single aggregated pricing node in Ontario, CA (ETIWANDA_06_N019).

Figure 7 shows a time-series of the price data that were collected. The chart also shows estimated energy savings for each hour. We were surprised by the fact that locational marginal prices did not vary much throughout the year. Regardless of the season marginal price on this node was usually 30-60 $/MWh. The price changed more as a function of the time of day than it did by season. Typically, the wholesale market clearing price is much larger in the summer months when loads are larger. Since dual-evaporative pre-cooling provides the most savings during the warmer months, the value of this strategy would be larger in a scenario where the variation in price was more pronounced and coincident with the savings.

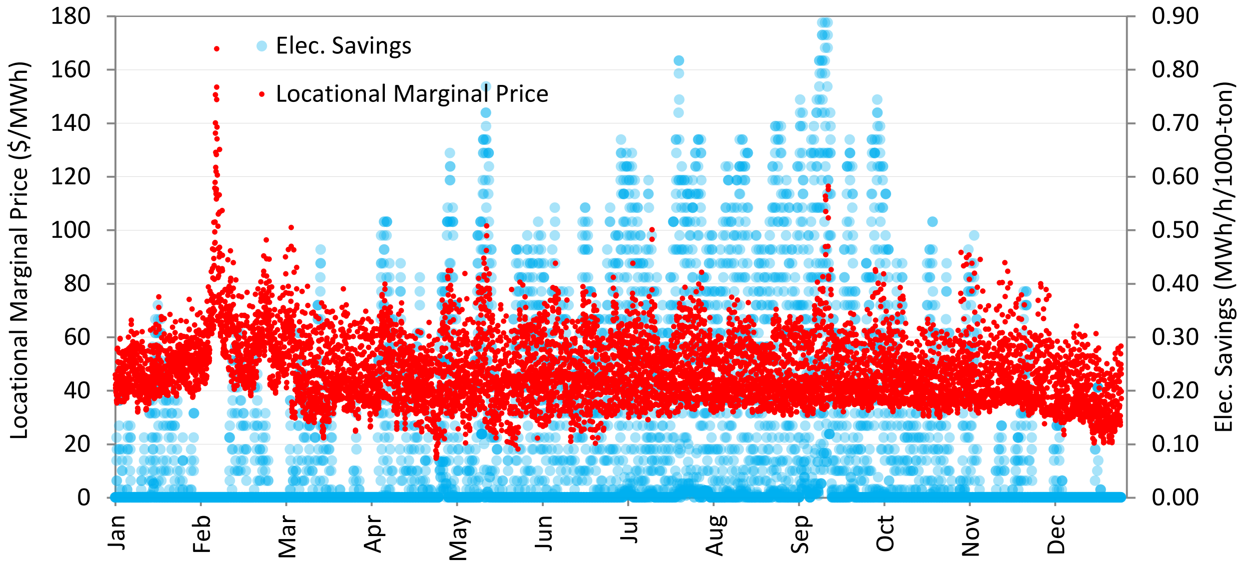

Figure 7: Time series trend for hourly electricity savings and the corresponding locational marginal price for the entire year simulated in Ontario California. (n=8760)

The capital costs and annual operating costs for the DualCool:

We obtained input from Integrated Comfort, the dual-evaporative pre-cooling manufacturer, to develop estimates about costs for the technology. For the results presented in this paper, we assigned a capital cost of $450 per ton. This represents current costs to a customer for capital and installation, and includes profit for the manufacturer and installing contractors. Some of the hypothetic business structures we consider would have the manufacturer provide all capital costs, and possibly maintain ownership of the equipment. In this case, the actual costs to the manufacturer could be lower than what we estimate here.

Annual costs for service and warranty were estimated to be about $14 per ton per year – this includes regular seasonal service visits, and the estimated cost of repair for any failures. The operating costs, which include web services and data, are about $4 per ton per year. In all of the business structures we consider, these annual costs are borne by the aggregator. Since it is in the aggregator's interest to monitor and upkeep the equipment in order to ensure performance, service and warranty would be provided for the lifetime of the equipment.

Four hypothetical business structures:

The model we developed calculates the value of energy savings, the reduction in demand charges, and the value of demand reduction compensated at the wholesale market clearing price. However, these financial benefits, and the costs required to facilitate the dispatchable efficiency measure can be applied in different ways. We explored the advantages and disadvantages of four different business structures that could deploy dual-evaporative pre-cooling for commercial air conditioners in order to fetch value on the wholesale spot market. We compare these models to the current practice, which does not seek demand reduction value on the wholesale market, and to other energy efficiency financing mechanisms.

In each business model, there are only two actors:

- The customer. A commercial building owner, or end user that purchases electricity from the utility at retail rates. The customer also pays demand charges for the maximum 15-minute-interval average power draw in each month.
- The efficiency aggregator. A third-party entity that monitors, controls, and aggregates hourly

demand savings for the dispatchable efficiency measure across a portfolio of commercial buildings. The aggregator may or may not own, service or maintain the dual-evaporative pre-cooler hardware, or the related air conditioning equipment, but they ultimately take responsibility for real time dispatch of the efficiency measure. A number of entities could serve this role, including the technology manufacturer, or utilities. An energy services company might also incorporate the strategy as one element in a more complete energy services portfolio.

Utilities, generators, and the wholesale market still exist, but their roles do not change for the business opportunities considered. Table 1 summarizes four business strategies, and the current practice - each option allocates costs and benefits to the two actors in different ways. The critical parameters are capital expenses for equipment, the costs for ongoing service, the benefit of reduced utility expenses, and the wholesale value for demand reduction. The way that each cost and benefit is allocated would dictate organization factors such as who owns what equipment, and whether or not the customer would pay the aggregator for ongoing benefits.

In every option we consider, the aggregator carries responsibility for seasonal maintenance, warranty service, remote fault detection, and real time control of the efficiency measure. In exchange, the aggregator receives all wholesale capacity value. In the current practice, the customer pays the manufacturer for all ongoing service activities - in fact, a multi-year service agreement and remote monitoring costs are required as part of the initial equipment purchase contract. In the four new business options all of these ongoing costs are borne by the aggregator.

The fact that ongoing service, warranty and maintenance are covered by a third party should serve as a significant motivator for new customers to adopt the measure. These services represent significant financial cost, and an organizational burden that many customers are hesitant to take on. The relationship also reduces investment risk for the customer. Since the aggregator commits to maintaining and controlling the systems for the purposes of real time grid capacity they can also guarantee savings for the customer in a reliable way.

One major advantage for each of these business options is that they do not introduce conflict with a customer's existing organizational relationships. The customer would continue on a regular billing structure with the utility, except they would use less electricity and incur lower demand charges. Existing equipment maintenance contracts, equipment warranties, and facility management contracts can remain in place. Existing energy management and control systems need not change, and the associated service contracts or organizational responsibilities can remain intact. This is appealing compared to some other energy services agreements that require significant organizational restructuring.

In Option 1 we explore the value proposition for each actor if the measure were installed at zero capital cost for the customer. This approach has some obvious advantages in terms of overcoming first cost hurdles. Additionally, it could allow the aggregator to maintain ownership of the dual-evaporative pre-cooling equipment, which might be appropriate if they are contractually responsible for maintaining its presence and operation. Third party ownership of the equipment could also be an effective way to bypass split incentive conflicts in non-owner occupied facilities.

Option 2 offers a scenario similar to the current practice - in that the customer pays all capital costs - but all ongoing service costs are borne by the aggregator in exchange for wholesale capacity value.

In Option 3 we consider an arrangement where the aggregator would guarantee a certain performance, then share in a portion of the energy and demand charge savings. This is effectively how many energy services companies operate already - in fact this measure could exist as one component in a larger energy services portfolio. In many energy services contracts the third party assumes responsibility for a customer’s utility bill, but sharing of the energy savings value could also be structured as an ongoing payment contract that allows the customer to maintain existing relationship with the utility.

With Option 4, we consider a possibility where the capital costs for the customer are substantially discounted, and the customer receives all value of utility bill savings for reduce energy use and demand. This option helps to assuage the first costs hurdle, and avoids the need for any ongoing financial agreements between the customer and the aggregator.

Table 1: Division of costs and benefits for five different business model options. In all options (excepting the current practice) the wholesale value of capacity from demand reduction is assigned to the aggregator.

	Costs		Benefits	
Business model option	Capital	Service, warranty, and controls	Retail energy and demand value	Wholesale capacity value
Current practice	Customer	Customer	Customer	N/A
Option 1: "Give it away"	Aggregator	Aggregator	Customer	Aggregator
Option 2: "Service and guarantee savings"	Customer	Aggregator	Customer	Aggregator
Option 3: "Split savings value"	Aggregator	Aggregator	Shared	Aggregator
Option 4: "Split capital costs"	Shared	Aggregator	Customer	Aggregator

These business models all rely critically on the ability to monitor and control the dual-evaporative pre-cooling equipment as dispatchable demand reduction across a portfolio of buildings - an arrangement that customers would have to agree to allow. At first glance it may seem that customers would not welcome the idea of third-party control over equipment on their premises. However, it is important to recognize that these controls would not interfere with the normal sequence of operations in a building - in fact they would really only be employed to temporarily disable the efficiency measure for short periods of time in order to establish a clear and reliable baseline. This sequence would be automated, it would not interfere with normal operations, and it would provide customers with reliable verification about the energy savings that are achieved.

The natural adoption of effective energy efficiency measures is stifled by a variety of factors that contribute to the so-called "efficiency gap". Overcoming these barriers requires well designed strategy.

> *"A successful energy efficiency [business] structure incentivizes each of the major stakeholders involved, and balances the relative risks of implementing energy efficiency improvements with the resulting energy savings returns and benefits." (Kim 2012)*

The proposed business models for dispatchable efficiency uniquely address most of the common barriers to adoption of energy efficiency, including:

- The first cost hurdle. The added value from providing demand reduction on the wholesale market enables the efficiency aggregator to partially or fully defray capital costs for the customer.
- Timing mismatch. Often the opportunity or interest to invest in energy efficiency does not align with facility operating cycles, and there is a tendency to avoid early retirement of inefficient equipment because of the financial disincentive to strand assets by devaluing equipment with remaining useful life. The dispatchable efficiency measure explored here is not limited by timing mismatch because it improves the efficiency of existing equipment, it can be deployed with value at any time, and it can be adapted and reused when baseline equipment is upgraded.
- Split incentives. The business structures presented here effectively bypass the split incentives problem since they can eliminate the need for capital expenditures by the customer, and since the physical assets deployed could be owned and operated by a third party.
- Scalability. The commercial building market is fragmented - motivations, constraints, and competing interests are different for every customer. We envision that the different business

models explored in this paper could be applied flexibly to suit the unique needs of each customer, without changing the overarching vision or requiring undue adaptation for application at scale.

- Existing property or financing restrictions. Since the dispatchable efficiency model can be structured to avoid capital expenditures, and would reduce a customer's monthly operating costs, the approach can effectively bypass restrictions that are associated with property financing.
- Energy baseline measurements. The dispatchable efficiency model is uniquely positioned to provide reliable baseline characterizations that evolve over time. This capability is an absolute must in order for demand reduction to be valued on the wholesale spot market, but it also improves the certainty of customer savings, and assuages risk associated with project financing.

RESULTS

Table 2 displays the value proposition for each of the five options. In all cases the customer's return on investment is outstanding. In the scenario modeled, the current business practice offers the customer a 200% return on investment and 2.75 year simple payback. All benefits in this scenario are accrued by the customer for retail energy and demand charge savings - the wholesale value of demand reduction is not included. In the baseline case the manufacturer and installer would be limited to one-time profits associated with the sale and installation of the equipment.

In Option 1, we find that the aggregator's income from demand reduction on the wholesale capacity market is not adequate to support "giving away" the capital investments and ongoing costs for service and controls. The aggregator's annual return on investment would be negative. This is unfortunate because the business strategy would help to bypass many of the barriers to adoption of energy efficiency. It should be noted that the marginal price for electricity during summer months might be much higher in other locations, in which case this business strategy would be more appealing. However, even if the business strategy were profitable for the aggregator, it would give away enormous value to the customer. The other options split costs and benefits more evenly.

Option 2 offers the customer a 360% return on investment and a simple payback of 2.5 years. The customer pays for capital, but the aggregator provides all service for no cost. In this case the aggregator also makes substantial profit - annual return is greater than 120%. Financially, this is nearly the best option for the aggregator. However, from the customer's point of view, this business approach may appear very similar to the current practices, and so customer recruitment might be stifled by the same barriers that currently limit adoption of energy efficiency – especially the first cost hurdle.

Option 3 presumes that the customer and aggregator can make an ongoing business arrangement to share value from the retail energy and demand charge savings – as might be facilitated by an energy services company. If the benefits from retail savings are split 50:50 the aggregator could cover all capital and ongoing service costs while capturing 127% return on investment. This approach is appealing because it removes the first cost hurdles, however the need for an ongoing financial agreement between the customer and aggregator is a major drawback.

Option 4 is advantageous because it reduces first cost for the customer, and avoids the need for ongoing financial arrangements between the aggregator and customer. For the scenario presented in Table 2, capital cost is split 50:50. This yields a mediocre return for the aggregator, and an excellent value proposition for the customer. Again, the wholesale price in other locations might make this business strategy more appealing for the aggregator, but for the location evaluated, splitting capital costs 50:50 doesn't seem like a worthwhile approach.

Figure 8 presents the return on investment for each party as a function of the fraction of capital paid by the aggregator. In every case the service and controls operation costs are borne by the aggregator, the benefits of retail savings go to the customer, and the value of wholesale capacity goes to the aggregator.

When all of the capital is paid by the aggregator, this is equivalent to Option 1, when none of the capital is paid by the aggregator this is equivalent to Option 2, and when the capital costs are split 50:50 this is equivalent to Option 4. The plot demonstrates that the business models might be tuned to divide the cost of capital in a way that is appealing to everyone involved. That said, the value proposition for the customer is excellent in every case. If a customer is not motivated to adopt for a simple payback of 2.5 years (as in Option 1) we are hesitant to expect that they would be much more likely to adopt for a payback of 1.25 years (as in Option 4). Again, the most strategic business structure might be one that provides a reasonable value proposition to the customer, but which does everything possible to ease the practical and organizational barriers to adoption.

Table 2: Return on Investment for Each Business Model

	Costs		Benefits		Return on Investment	
Business model option	Capital	Service, warranty, and controls	Retail energy & demand value	Wholesale capacity value	Aggregator	Customer
Current	Customer	Customer	Customer	N/A	N/A	206%
Option 1	Aggregator	Aggregator	Customer	Aggregator	(26%)	∞
Option 2	Customer	Aggregator	Customer	Aggregator	121%	360%
Option 3	Aggregator	Aggregator	Shared 50:50	Aggregator	127%	∞
Option 4	Shared 50:50	Aggregator	Customer	Aggregator	11%	820%

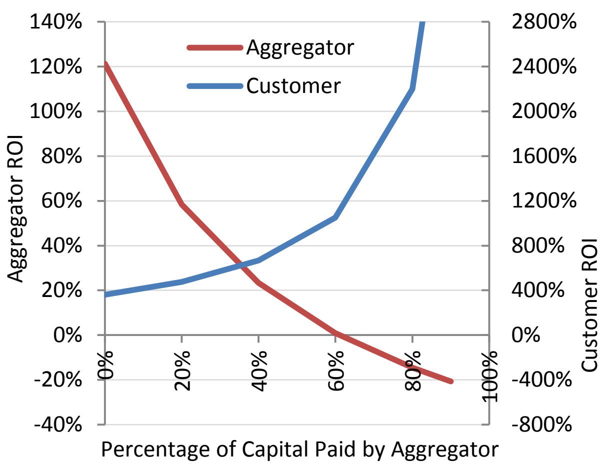

Figure 8: Return on investment for each party in Option 4 as a function of the fraction of capital paid by the aggregator

In this light, regardless of the way that financial costs and benefits are divided, an aggregator ought to emphasize the non-energy advantages of the dispatchable efficiency arrangement including:

1. Guaranteed energy and demand savings
2. No cost service
3. Extended equipment longevity
4. Bypass principal agent problem
5. Extend life of existing equipment
6. Simpler than ESCO agreements
7. Does not disturb existing structures
8. Can retrofit whole building without scrapping existing systems

Discussion and Conclusions

A variety of building energy efficiency measures could be controlled in real time to respond to the changes in net load on the grid. These dispatchable efficiency measures could operate to reduce energy use at any point in time; but operation during peak demand periods would have the most benefit in terms of avoided wholesale electricity costs. FERC Order 745 directs that demand reduction resources should be compensated by the wholesale electricity market at the locational marginal price – akin to conventional generation resources. In this paper we explore several business strategies for an aggregator that could seek value for the demand reduction generated by dual-evaporative pre-cooling for commercial air conditioners across a regional portfolio of buildings.

We find that the wholesale market value that could be gained from this measure as a demand side resource is substantial. In fact, this financial value creates business opportunities that could help to relieve some of the qualitative barriers to adoption of energy efficiency. Some business approaches could install the measure a no cost to the customer and still make a substantial rate of return. All of the business models considered provide ongoing service, and warranty for the efficiency measure, at no cost.

The results presented in this paper are based on one year of hourly day ahead locational marginal prices for one aggregated pricing node in Ontario California, the coincident weather conditions for that region, reasonable inputs about equipment costs and retail energy costs, and a simple model for thermal load and air conditioner efficiency as a function of outside temperature. The results will be different for other scenarios.

Most importantly, the wholesale price in Ontario was not as sensitive to season and temperature as would be expected for the grid as a whole. Therefore, the same exercise for other regions would likely generate more appealing financial opportunity for an efficiency aggregator that seeks value from the wholesale market.

We recommend that future efforts should explore the way that these business options perform in different regions and building applications. Demand reduction from this measure is most significant during peak demand periods, which should correspond with the highest market clearing prices. Future efforts should also be careful to account for rules about current practice for how demand side resources are compensated. For example, FERC 745 directs that demand side resources should only be compensated when paying for demand reduction at the market clearing price produces net benefits. We argue that demand reduction provides benefit at any point in time, and that it should be compensated at all times at a rate that is equal to the benefits it provides. In our simplified model we pay the locational marginal price for demand reduction at any time, but this does not align with current market practice.

REFERENCES

Borlick, Robert. “Paying for Demand-Side Response at the Wholesale Level: The Small Consumers’ Perspective.” The Electricity Journal 24, no. 9 (November 2011): 8–19. doi:10.1016/j.tej.2011.10.012.

California ISO. “Capacity Markets Background Information,” n.d. https://www.caiso.com/Documents/CapacityMarketGeneralOverview.pdf.

California ISO Open Access Same-Time Information System (OASIS), n.d. http://oasis.caiso.com/.

Efficiency Vermont. “Efficiency Vermont Annual Report 2014.” Efficiency Vermont, October 7, 2015. https://www.efficiencyvermont.com/docs/about_efficiency_vermont/annual_reports/evt-annual-report-2014.pdf.

Falk, Jonathan. “Paying for Demand-Side Response at the Wholesale Level.” The Electricity Journal 23, no. 9 (November 2010): 13–18. doi:10.1016/j.tej.2010.10.001.

Hogan, William W. “Demand Response Compensation, Net Benefits and Cost Allocation: Comments.” The Electricity Journal 23, no. 9 (November 2010): 19–24. doi:10.1016/j.tej.2010.10.002.

Kim, Charlotte, Robert O’Connor, and Kendall Boden. “Innovations and Opportunities in Energy Efficiency Finance.” Wilson Sonsini Goodrich & Rosati, May 2012. https://www.wsgr.com/publications/PDFSearch/WSGR-EE-Finance-White-Paper.pdf.

Woolley, Jonathan. “Western Cooling Challenge Laboratory Results: Trane Voyager DC Hybrid Rooftop Unit.” University of California, Davis, October 31, 2012. http://wcec.ucdavis.edu/wp-content/uploads/2014/01/TraneVoyagerDC_LaboratoryResults-Repor.pdf.

Woolley, Jonathan, and Caton Mande. “Performance Evaluation for Hybrid Rooftop Air Conditioners with Dual Evaporative Pre-Cooling.” Southern California Edison, October 2014.

Xu, Lin. “Demand Response Net Benefits Test.” California ISO, June 2011. https://www.caiso.com/Documents/IssuePaper_DemandResponseNetBenefitsTest.pdf.