Using Insurance to Manage Reliability in a Distributed Electricity Sector: Insights from an Agent Based Model

Rolando Fuentes[1]* and Abhijit Segupta[2]

[1] Research Fellow, King Abdullah Petroleum Studies and Research Center (KAPSARC). Airport road, Riyadh, Saudi Arabia. *Corresponding author: rolando.fuentes@kapsarc.org
[2] Senior Lecturer in Innovation and Entrepreneurship. Kent Business School | University of Kent, Canterbury, United Kingdom

Abstract

High penetration of new distributed energy technologies would call for a different way to manage reliability in the power sector. In the most extreme case, the consumer would produce all the power their home needs and store it until the time of consumption, effectively by-passing the utility. The issue of electricity security for self-sufficient households would still remain because these technologies can eventually fail due to technical or weather conditions, or be subject to spikes in demand which installed capacity will not be able to meet. An option could be to draw power from the grid when this happen, but this behaviour could eventually constitute an existential threat for utilities. In response to these concerns, we test in this paper the creation of a risk market that enables reliability preferences to be internalized through the use of insurances. We propose the utility can offer last resort power – an insurance-- to energy self-sufficient households to protect them against the prospect of a blackout. The overarching idea is that instead of selling commoditized kilowatt-hours, consumers would pay for guaranteed services.

Introduction

High penetration of new distributed energy resource (DER) technologies would call for a different way to manage reliability in the power sector. In central electricity systems, planners invest in capacity higher than the peak load to have reserve margins that can deliver almost perfect coverage, as they can spread this cost among all customers. This approach was more feasible when technology options were limited, and the underlying assumption was that blackouts have an infinite negative value. In a DER-dominated power system, however, it seems unnecessarily expensive and unfair, as all customers pay an equal amount regardless of their preferences for risk. To deal with security of supply in this context, this paper tests the creation of a risk market that enables reliability preferences to be internalized through the use of insurance.

New distributed energy technologies - by this we refer to the combination of domestic photovoltaic (PV) panels plus batteries plus information devices - allow households to generate, trade, reduce, and shift their electricity consumption, bypassing the traditional utilities. The most extreme case of distributed energy resource (DER) adoption would be self-sufficiency, i.e., the consumer produces all the power their home needs and also stores enough for later use. However, the issue of security of supply for self-sufficient households would remain. There is an inherent risk in all intermittent technologies that they could temporarily fail due to adverse weather or technical problems.

Being self-sufficient would allow households to disconnect from the grid and avoid all external charges (Green and Staffell 2017). However, most of those households will still need to use the networks to draw power from the grid if their system fails. A pay-as-you-go scheme would not reflect the opportunity cost of idle infrastructure and could lead to a utility death spiral if the way this back up is priced stays the same (Felder and Athawale 2014; Muaafa et al. 2017). Traditionally, network costs are bundled into the price of electricity. This pricing may be inadequate if self-sufficient households defect from the grid on mass, and in the presence of technologies where customers can best reflect their preferences for guaranteed services.

Insurance exists to reduce or eliminate the cost to an individual who faces an adverse event, like the loss of power. There are many different types of insurance policies available, and virtually any individual or business can find an insurance company willing to insure them for a price. We propose the utility can offer power of last resort to energy self-sufficient households in case they suffer a blackout. Specifically, the value proposition is that utilities' idle capacity resulting from the fast deployment of DERs in the domestic sector can be repackaged and repriced as insurance. This service is currently embedded in electricity provision and is taken for granted. By selling insurance, utilities would guarantee a stable revenue stream by charging a constant fee, instead of charging very high prices in periods of abrupt demand surges followed by low fees for periods of negligible demand. This proposal may also appeal to customers as an insurance-based model can reflect their attitudes towards the risk of blackout and can, therefore, be an efficient and equitable way to pay for network access.

We investigate the extent to which contracts for insurance can converge into a theoretical optimal contract. This contract is defined on the basis of a household-specific energy budget, a risk aversion attitude, an expected loss and with the supplier possessing full and complete information about the households. The static version of the model is then generalized into a dynamic framework where households are allowed to renew or switch contracts over time, and firms do not possess full information about households and hence cannot offer an optimal contract. Instead, they offer a menu of contracts from which households choose the one they prefer and can most afford. Firms are allowed to adapt their offerings periodically based on profitability considerations.

The dynamic outcomes of this generalized setup are examined through an agent-based model, particularly how the perception of risks and losses impact the price of insurance and potential revenues of utilities. We also use the agent-based model to examine the stability of such an insurance market and household-level outcomes, and explore how this market would react to policy initiatives that affect its underlying parameters.

This paper addresses a latent problem. While most of the literature on disrupting DERs looks at the gradual penetration of new technologies, we take a different stance and assume that this penetration has already occurred. So, instead of analyzing the incumbent utilities' reaction to the erosion of their market shares, we start from the end and figure out the steps back.

Through this paper's proposed approach, the security of electricity provision would move away from an engineering approach, based on system-wide costs, towards an economic approach, where an insurance mechanism along with an array of alternative contracts can balance system-wide supply and demand.

Electricity has multiple attributes beyond the energy component. These additional dimensions may, at times, be difficult to price. For example, non-priced factors such as levels of emissions attached, reliability services and supply risks are embedded into the price of electricity. These dimensions of electricity can be treated as externalities as they are not explicitly captured in the pricing systems. As markets for these attributes do not exist (for example, customers often take reliability for granted), there is a temptation to think that the un-marketed good is un-marketed because it is abundant and hence has a negligible price (Pearce 2000). As such, consumers cannot reflect their preferences for reliability.

The only proposal for an insurance market for electricity has been from the head of the United Kingdom's energy regulator, the Office of Gas and Electricity Markets, or OFGEM, in an article published in The Telegraph titled "Households could be charged annual 'insurance premium' for access to electricity grid." (Gosden 2016) Other business models that unbundle electricity services have started to emerge. For example, the Rocky Mountain Institute proposed a business model for lighting — measured in lumens — where the provider delivers a specified service (Calhoun et al. 2017). Other institutes such as Energy Systems Catapult have applied broadly similar thinking to propose the creation of a domestic heating and cooling service (Watkins 2017).

Model

The following assumptions for our model provide the context for our analysis (see Fuentes, Blazquez and Adjali [2019] for the justification of these shortcuts):

1. Households wish to have power provision 100 percent of the time, at the lowest cost to them, while maximizing their level of electricity independence.
2. Overnight, all households install large amounts of PV and batteries, which allows them to be power independent.
3. Utilities are entitled to cut households off from their network if they no longer use their service on a regular basis.
4. There is an inherent risk associated with all intermittent forms of generation.
5. Households cannot buy power from other households. This simplification allows us to isolate pure risk management strategies and ignore potential externalities arising through inter-household networks.

6. Households have pre-allocated budgets for purchasing additional energy or insurance in the event a shortfall occurs from their domestic source.
7. For further simplification, we do not analyze commercial and industrial users.

Theoretical foundations

This section provides the theoretical basis upon which we test the prospects of an insurance market for electricity. Please refer to the appendix for a technical description of this model.

We first, based on a static framework, derive the optimal conditions for a household to buy an insurance contract for electricity given a budget constraint, its risk preferences and expected power loss. The risk to the household arises from the fact that its total power demand may exceed its installed capacity with a certain degree of probability at any given time. The budget constraint is defined as a fixed pre-allocated budget for additional energy needs for each household.[3] Using a simple single period model, we find households would be willing to purchase a contract if the degree of cover offered (κ) is sufficiently high (a threshold defined by the household specific loss expected as well as its own risk profile).

Next we identify what the 'ideal' contract is for a household from the supplier's perspective. We find that, in equilibrium, the supplier can maximize its profit from the household by offering it a two-part tariff contract with a per unit price of energy consumed and a 'standing charge' price. These are functions of the household's budget, expected loss amount and risk profile, as well as the supplier's fixed costs.

We extend this static model into an inter-temporal dynamic model and relax the rationality assumptions. We also introduce an alternative dynamic specification, where both the supplier and the households are allowed to learn from past behavior and outcomes and adapt their behavior over time. This allows households to make 'mistakes' in their choice of insurance contract (coverage purchased) and the ability of the utility to change the tariff for each contract offered based on demand.

These adjustments allow us to incorporate learned behavior. First, we introduce a coverage limit in the contract, which acts as an additional constraint for households in the model. Second, we

[3] See appendix for more details.

allow for the possibility of the supplier making multiple insurance contracts available to households. In this scenario, the supplier provides a menu of alternative contracts, from which households have to select one. Third, we introduce an adaptation rate parameter τ, which represents how quickly households and the supplier are allowed to change their respective demand and supply behaviors. This makes the modeling less restrictive and provides us with a more flexible and dynamic setup, the implications of which we test using an agent-based model. The resulting behaviour is later benchmarked against the rational equilibrium outcome, obtained from the static model.

The sequence of the dynamic model is as follows. At time $t = 0$, the supplier offers the contract that bundles specific prices and expected loss coverage, from which each household selects one based on its risk profile, energy budget and required coverage. In each subsequent period the households may face a random loss of energy, which may or may not be covered in full, based on the contract purchased and the coverage it offers. The supplier updates the contract every τ periods, and the households, once they have chosen a contract, are contractually bound to it for τ periods. After every τ periods, the supplier has the option of modifying the offerings and households have the option to update their choices. This set of contracts and choices once again remain fixed for the next τ periods.

A household n selects contract k with probability ρ_t^k at every τ periods from a menu of alternative contracts. The probability distribution is updated in every period t, thus making ρ_t^k a household-specific endogenous parameter. The update rule is based on learning from past 'mistakes' that arise from choosing a contract with an inappropriate level of coverage, due to their budget constraint, risk profile, or exogenously fixed coverage limits. Hence, at the time of choosing the next contract, the household faces a potentially new probability distribution, representing the cumulative effect of the mistakes from the last τ periods. The supplier is allowed to adjust the *unit prices offered* in the contracts every τ periods, reflecting changing demand conditions. The existing per unit price is adjusted (increased or decreased) to reflect the change in cumulative demand over τ periods, provided the change is large enough, as determined by parameter θ. A higher θ implies a *less frequent* price update, but each update is of a greater amount than a smaller θ; a smaller θ implies more frequent updates but by a smaller proportion than a higher θ.

Agent-based model

The dynamic model described above is operationalized as an agent-based model. The model solves for final unit prices under different contract settings, cumulative revenues for the supplier, and the extent to which this product allows reliability to be internalized by calculating the percentage of houses that are left uninsured.

Each household agent is characterized according to its risk aversion, expected losses, budget and a probability of distribution, as shown in Table 1. The supplier provides a menu of three alternative contracts for households to choose from, i.e., $K = 3$. A household at $t = 0$ starts with an even probability distribution (0.33, 0.33, 0.33) across all contracts in the menu. The households choose a contract and the menu is updated every 12 time steps within the simulation ($\tau = 12$). By default, we set contract 1 to be a 'spot' contract with no pre-specified energy limit, i.e., potentially $L^1 \rightarrow \infty$. Contracts 2 and 3 set upper limits on how much energy a household can draw. We impose the following constraints on the contracts: $L^1 > L^2 > L^3$ and $p^1 > p^2 > p^3$. Thus in a sense, the spot contract is the risk-less contract but with the highest per unit price, while contract 3 is the most risky contract but with the lowest per unit price. Contract 2 is an intermediate one, which provides a balance between risk and cost.

Table 1. List of parameters and operational values in the simulations

Parameter	Context	Meaning	Values
Input			
$\bar{\alpha}$	Household	Upper limit of household-specific coefficient of absolute risk aversion	{0.1, 0.9}
$\bar{\bar{L}}$	Household	Upper limit of potential loss per household	{100, 500}
$\bar{B}$	Household	Upper limit of household-specific budget	{50, 500}
$\bar{\pi}$	Household	Upper limit of household-specific probability of loss	{0.2, 0.8}
p_0^1	Supplier	Initial unit price for contract 1 (at $t = 0$)	{100, 150}
L^1	Supplier	Limit on energy that can be drawn under contract 1, as a proportion of upper limit of potential loss ($L^1 * \bar{\bar{L}}$)	{1}
p_0^2	Supplier	Initial unit price for contract 2 (at $t = 0$)	{25, 75}
L^2	Supplier	Limit on energy that can be drawn under contract 2, as a proportion of upper limit of potential loss ($L^2 * \bar{\bar{L}}$)	{0.5, 0.9}
p_0^3	Supplier	Initial unit price for contract 3 (at $t = 0$)	{1, 10}
L^3	Supplier	Limit on energy that can be drawn under contract 3, as a proportion of upper limit of potential loss ($L^3 * \bar{\bar{L}}$)	{0.1, 0.45}

θ	Supplier	Threshold in cumulative demand change for price adjustment	{0.05, 0.5}
Control			
$\bar{U}$	Household	Utility under certain zero loss	{2500}
N	Environment	Number of households	{2500}
F^1	Supplier	Lump sum payment under contract 1	{0}
F^2	Supplier	Lump sum payment under contract 2	{50}
F^3	Supplier	Lump sum payment under contract 3	{50}
Output			
p^1_{300}	Supplier	Final unit price under contract 1	
p^2_{300}	Supplier	Final unit price under contract 2	
p^3_{300}	Supplier	Final unit price under contract 3	
Revenue	Supplier	Cumulative revenue for τ previous periods	
h^F	Environment	Houses (% of those facing a loss) with full coverage, averaged across 300 runs	
h^B	Environment	Houses (% of those facing a loss) with partial coverage due to binding budget constraint	
h^L	Environment	Houses (% of those facing a loss) with partial coverage due to binding contract coverage limit constraint	
h^I	Environment	Houses (% of those facing a loss) with partial coverage due to binding incentive constraint	
h^{BL}	Environment	Houses (% of those facing a loss) with partial coverage due to binding budget constraint and contract limit constraint	
$avg\ \rho^1, sd\ \rho^1$	Household	Mean and SD of the distribution of $\rho^1(300)$ across households	
$avg\ \rho^2, sd\ \rho^2$	Household	Mean and SD of the distribution of $\rho^2(300)$ across households	
$avg\ \rho^3, sd\ \rho^3$	Household	Mean and SD of the distribution of $\rho^3(300)$ across households	

The update of a household's probability distribution over the available menu of contracts assumes a simple reinforcement learning algorithm, based on which of the three constraints had been binding in period $t-1$. These are summarized as follows:

- Rule 1: If the budget constraint was binding in period t, in period $\mathrm{t}+1$ reduce the probability weights on contracts with relatively higher per unit prices and increase the probability weights on contracts with relatively lower per unit prices.
- Rule 2: If the incentive constraint was binding in period t, in period $\mathrm{t}+1$ reduce the probability weights on contracts with relatively lower coverage limits and increase probability weights on contracts with relatively higher coverage limits.

- Rule 3: If the coverage limit constraint was binding in period t, in period $t+1$ reduce the probability weights on contracts with relatively lower coverage limits and increase probability weights on contracts with the highest coverage limits.

Rules 2 and 3 are similar as they both add weight on contracts with relatively higher coverage. The only distinction is in the rate at which they shift the households' choice towards the spot contract, with Rule 3 shifting the probability weights in favor of the spot contract at a higher rate than Rule 2. Rule 1, on the other hand, shifts the weights in favor of cheaper contracts. Once the updates have taken place, the probabilities are rebased to ensure that they are bounded within the [0, 1] interval.

Experimental setup

The agent-based simulations are implemented in Netlogo 6. Each run of the simulation represents one instance of a single experiment, where an experiment is defined as a unique combination of input parameter values. Each experiment is replicated 10 times to account for randomness in the model, where each replication is labeled as a 'run.' Each run is composed of 300 time steps because we observed that the model attained a high degree of stability well within this figure.

The input and output variables were recorded over multiple text files by Netlogo at every time step of each run in the experiment. The text files were then combined and analyzed using the statistical software R separately. The analysis of the results is presented in the three subsections of the results.

Results

The model solves for final prices for each one of the three type of contracts, the revenue for utilities, and the extent to which households internalize and transfer risks to the utility. We measure these by estimating the percentage of houses that obtain full or partial coverage. For those who only obtain partial coverage, we investigate the source of that decision, i.e., whether the decision was constrained by their budget, by contract limits and/or binding incentives.

To give a comprehensive picture, we present the convergence and stability properties of the model based on the central values, variances and trends in output variables, measured over the 300 steps in each run. We then explain the observed variation in the outputs at the macro level, using multivariate regressions, and present the estimated impact of input variables. We measure the values of output variables at the end of each run in each experiment, or in other words the 'final' value of the output at the 300th step. Finally, we examine sub-samples of households and correlate their behavior with their intrinsic properties to explore the behavior at the agent level. This is done using a final and intermediate measure of outputs in each run.

Convergence

We find that the model has strong convergence properties in the output variables under all input settings explored within the simulations. This is true for prices, the proportion of houses facing one or more (or no) constraints, and in the distribution of choice probabilities. This indicates that, conditional on the learning and adaptation algorithms of the supplier and households, the market moves towards a dynamically stable outcome within the stipulated 300 steps.

Figure 2. Convergence patterns in the unit prices of the three contracts

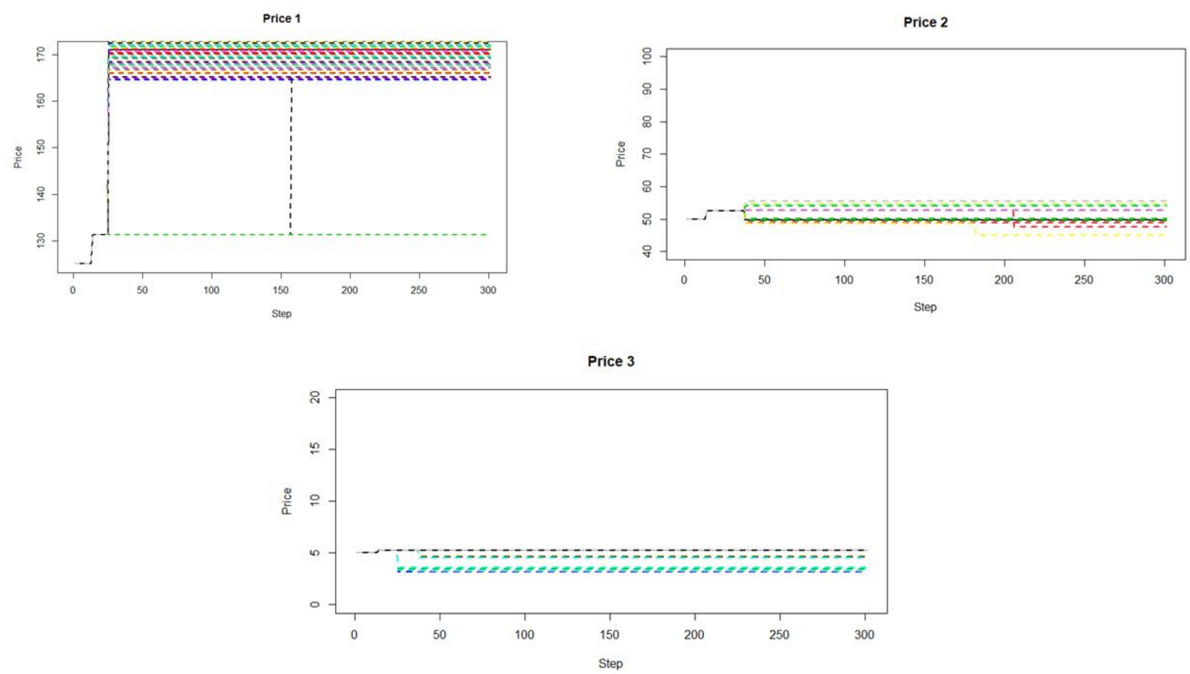

Figure 3. Convergence patterns in the average choice probabilities across the three contracts

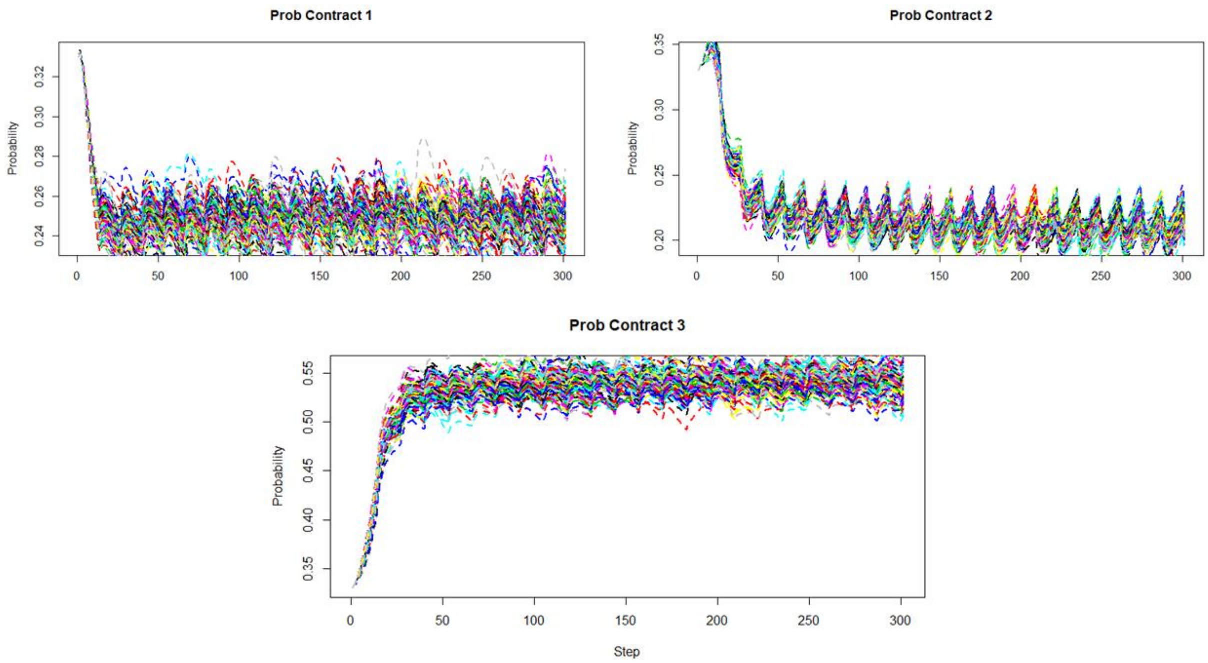

Figure 4. Convergence patterns in the proportions of households facing alternative constraints

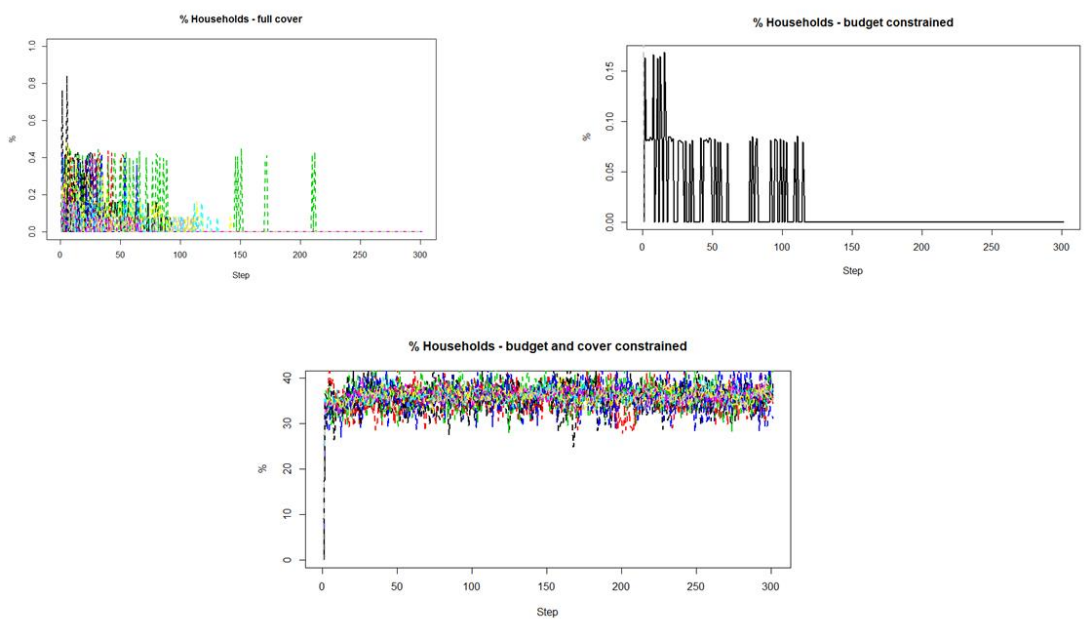

Figures 1-3 show outputs of a sample of runs from the sensitivity analysis, where the control variables are varied across two levels – low and high. What we observe is that the output variables either achieve a stable value quickly (such as in the prices in Figure 1), or the variations exhibit a predictable pattern within a very narrow band (such as in choice probabilities and household proportions in Figures 2 and 3, respectively). These patterns are consistent in the actual experiments, where input variables are varied and controls are held fixed.

This implies that the dynamic model implemented here exhibits strong convergence properties, even without the strict rationality assumptions of the static case. A simple reinforcement learning mechanism can provide stable outcomes in the market.

Input-output relationships

Here we present the partial impact of input variables on output variables using a set of linear regressions, in which the coefficients of the input variables indicate the *partial* impact of the input on the output. Tables 2-4 show the magnitude of the impact, whether the impact is positive or negative and the statistical significance of independently varying each input on the output, as measured in the experiments. The values presented in bold indicate a relatively large magnitude of impact (which is statistically significant) when compared with the size of other coefficients in the same regression equation.

Impact on contract prices and revenue

We observe that as the potential loss for households increase, they generally switch from cheaper but riskier to more expensive but safer contracts. The unit prices of the less risky contracts (1 and 2) seem to depend positively on the upper limit of expected losses faced by the households, as seen by the coefficients of $\bar{\pi}$. However, the price of contract 3, the riskiest contract, has a negative and significant coefficient. Thus, higher levels of expected loss faced by households pushes up the demand for the safer contracts with relatively wider coverage, with households willing to pay the higher unit price.

Table 2. Effect sizes on output variables – final prices and revenue

Output / Input	p^1_{300}	p^2_{300}	p^3_{300}	revenue
Intercept	-44.18***	-12.91***	2.05***	3.27E04***
$\bar{\pi}$	**28.36*****	**4.20*****	**-0.80*****	**7.48E04*****
$\bar{\bar{L}}$	0.00	0.00	-0.0001***	23.39***
$\bar{\alpha}$	-0.06	-0.22	-0.02	-45.64
$\bar{B}$	-0.02***	-0.01***	0.00	247.70***

p_0^1	1.37***	0.00	0.00	3.58
L^2	**83.23*****	**12.26*****	**-2.45*****	**2.27E04*****
p_0^2	-0.003	1.10***	0.00	9.98
L^3	**59.64*****	**22.60*****	**-2.31*****	**7.46E04*****
p_0^3	0.38***	0.08***	**0.93*****	**1.55E03*****
θ	**-160.00*****	**-11.60*****	**2.70*****	**3.22E03***

We observe that the demand for the riskiest contract not only falls, as expected, when its nearest rival becomes less risky, but it also falls when it becomes less risky. Less frequent but larger price changes are associated with increased demand for less risky alternatives and lower demand for the riskier contract. Overall, it seems that the demand for the less risky contracts is more sensitive to changes in the above input parameters than the riskier ones. This is apparent in the gradual decrease in the size of the coefficients from a high positive value (contract 1) to a small negative value (contract 3).

The utility's revenue is also most strongly impacted by the same parameters as above, plus the initial price of contract 3, p_0^3. Of all the initial prices, only the initial price of contract 3, the riskiest contract, has a notable impact on its final price, affecting it positively and, in turn, positively impacting the utility's revenue. This is another way in which the riskiest contract is distinct from the other two.

Constrained households

One would expect that increasing coverage would reduce the proportion of households constrained by coverage limits, and possibly increase the proportion obtaining full coverage. However, as shown in Table 3, this is not always the case (negative coefficient of L^3 for h^F, and a positive coefficient of L^2 for h^L). We also see that varying L^2 and L^3 also seems to affect the proportion of households subject to the budget constraint, which is a surprising result. No households under the settings used in our experiments faced a binding incentive constraint (where the actual loss was greater than the certainty equivalent).

Table 3. Effect sizes output variables – household coverage and constraints

Output	h^F	h^B	h^L	h^{BL}	h^I

Input					
Intercept	6.79***	44.65***	6.79***	48.07***	NA
$\bar{\pi}$	-0.4***	0.18	-0.49***	0.63***	NA
$\bar{\bar{L}}$	-0.01***	0.01***	-0.01***	0.01***	NA
$\bar{\alpha}$	-0.01	-0.02	-0.01	0.05	NA
$\bar{B}$	0.02***	-0.02***	0.02***	-0.03***	NA
p_0^1	0.00	0.00	0.00	0.00	NA
L^2	**5.49******	**9.47******	**5.48******	**-15.54******	NA
p_0^2	0.00	0.01***	0.00	0.00	NA
L^3	**-17.72******	**34.53******	**-17.72******	**-27.31******	NA
p_0^3	-0.6***	0.32***	-0.67***	0.67***	NA
θ	**-1.23******	**0.85******	**-1.23******	**1.14******	NA

Choice probabilities

As losses become more likely, the probability of choosing contract 2 increases. This seems to suggest that the increased chance of households facing a loss leads to an increased preference for the middle contract, which is both relatively less risky than contract 3 and less expensive than contract 1. There is a still a spread around the choice probabilities across households.

The results presented in Table 4 suggest that the upper limit of the household-specific loss ($\bar{\pi}$) and the coverage limits (L^2 and L^3) are the only input parameters which have a notable impact. Thus there is an increased preference for the 'safer' option.

Table 4. Effect sizes output variables – choice probability means and standard deviations

Output / Input	$avg\ \rho^1$	$sd\ \rho^1$	$avg\ \rho^2$	$sd\ \rho^2$	$avg\ \rho^3$	$sd\ \rho^3$
Intercept	0.49***	0.44***	0.20***	0.18***	0.31***	0.37***
$\bar{\pi}$	**-0.06******	**-0.07******	**0.11******	**0.04******	**-0.04******	**-0.03******
$\bar{\bar{L}}$	0.00	0.00	0.00	0.00	0.00	0.00
$\bar{\alpha}$	0.00	0.00	0.00	0.00	0.00	0.00
$\bar{B}$	0.00	0.00	0.00	0.00	0.00	0.00
p_0^1	0.00	0.00	0.00	0.00	0.00	0.00

L^2	**-0.18*****	**-0.12*****	0.00	0.00	**0.17*****	**-0.02*****
p_0^2	0.00	0.00	0.00	0.00	0.00	0.00
L^3	**-0.24*****	**-0.03*****	**-0.20*****	**-0.02*****	**0.44*****	**0.09*****
p_0^3	0.00	0.00	0.00	0.00	0.00	0.00
θ	0.00	0.00	0.00	0.00	**0.01*****	0.00

Interactions of inputs

Figures 4-7 show how output variables are affected by each input for the *entire range* of values of the other input variables. We then concentrate on the joint impact of the four primary inputs - $\bar{\pi}$, L^2, L^3 and the threshold parameter θ. [4]

Price variation per contract

Figures 4a, b and c explore the variation in the final unit prices in the three contracts across different input parameter settings. As can be seen, θ is a strong moderator of the relationships between the rest of the inputs and prices. We find there are consistent relationships between the unit prices and $\bar{\pi}$, L^2, L^3 at low values of θ. However, these relationships disappear almost completely for higher values of θ. Therefore, in situations where the utility makes infrequent but large changes in prices, outputs no longer depend on the other parameters. This is an important result in terms of market design and has important regulatory implications.

[4] Figures 4 to 7 graphically present the impact of the four inputs on every output variable. In each figure, each box and whiskers plot shows the four quartiles and the median of the output variable in question across all the 10 runs in any given experiment. Each figure is divided into eight regions, with two boxplots in each region. These two boxplots represent experiments with low (0.2) and high (0.8) values of $\bar{\pi}$ respectively for given values of the other three inputs. The left four regions represent experiments with the lower value of θ=0.05, whereas the four regions on the right half of the figures represent the higher value of θ=0.5. Each half is further sub-divided into two quarters – where each quarter is representative of low L^3 (0.1) and high L^3 (0.45) respectively. Finally, each quarter is further sub-divided into two sections, where each section represents experiments with the low value of L^2 (0.5) and the high value of L^2 (0.9). As mentioned above, there are eight such sections, with two boxplots in each. The ordering of experiments in each figure is consistent across Figures 4-7, to make them comparable.

Figure 5a. Variation in unit price of Contract 1

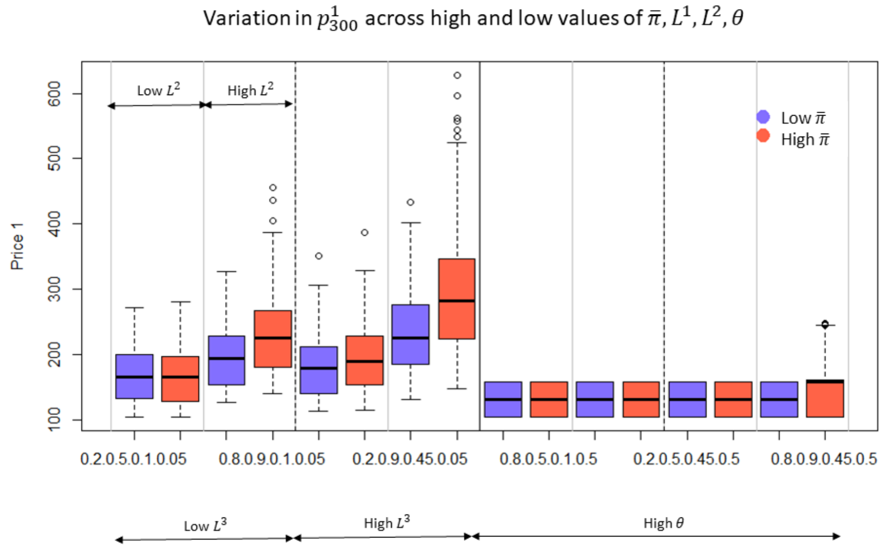

Figure 5b. Variation in unit price of Contract 2

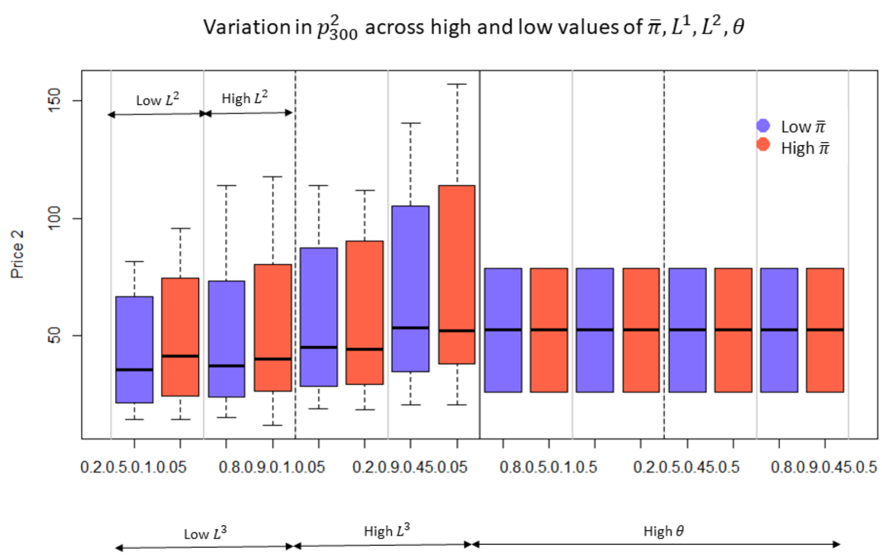

Figure 5c. Variation in unit price of Contract 3

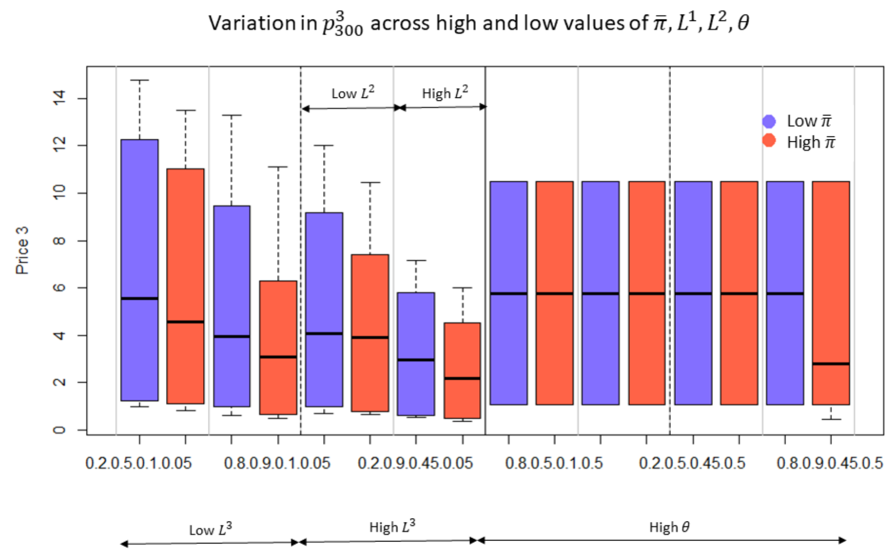

We noted other significant interaction effects between the input and output variables, such as the positive $\bar{\pi}\ x\ L^2$ effect on ρ^1, implying that the fall in ρ^1 due to increased $\bar{\pi}$ is steeper under high values of L^2. Another significant interaction can be seen for ρ^2, where $\bar{\pi}\ x\ L^3$ is negative, and the negative $\bar{\pi}\ x\ L^2$ effect on ρ^3.

Figure 6. Variation in supplier revenue

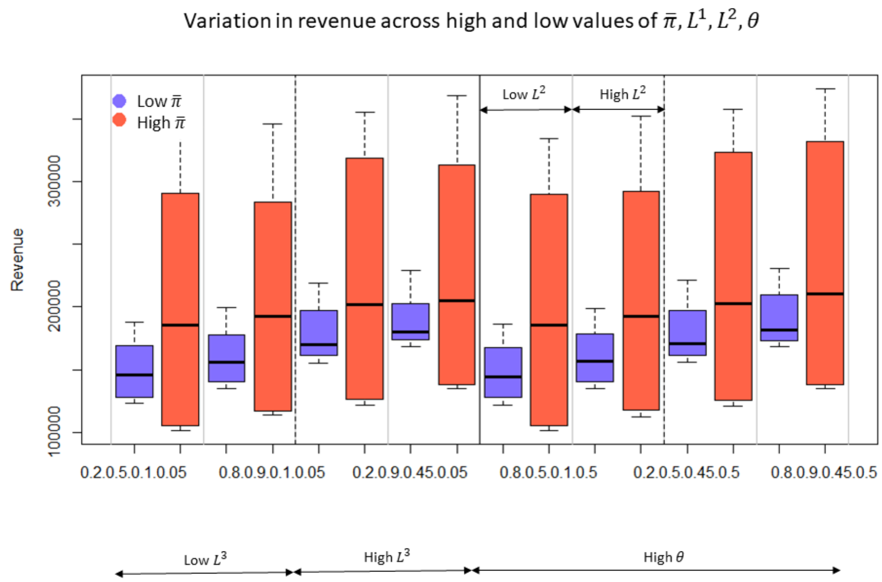

Figure 7a. Variation in average probability of choice of contract 1

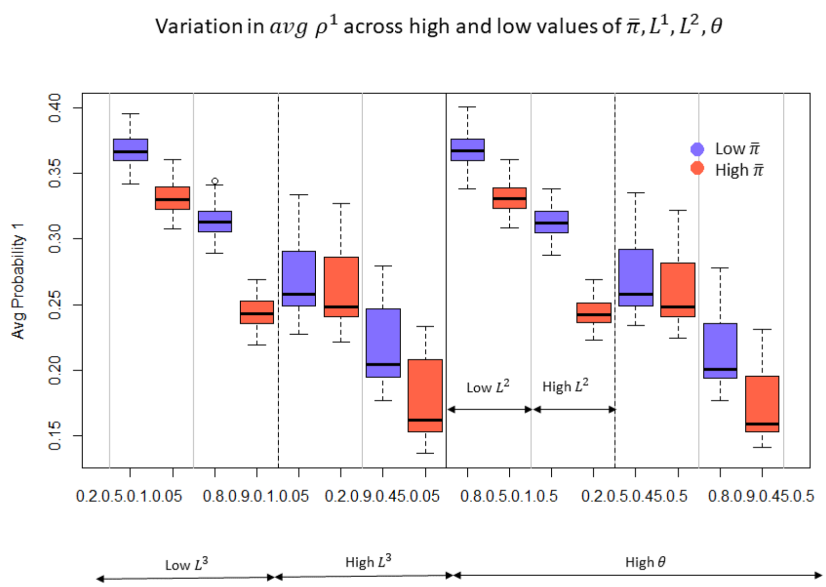

Figure 7b. Variation in average probability of choice of contract 2

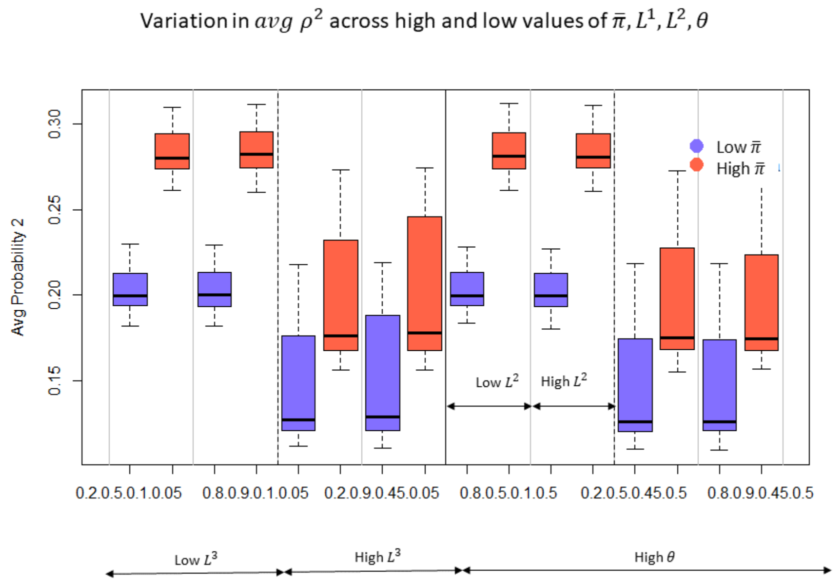

Figure 7c. Variation in average probability of choice of contract 3

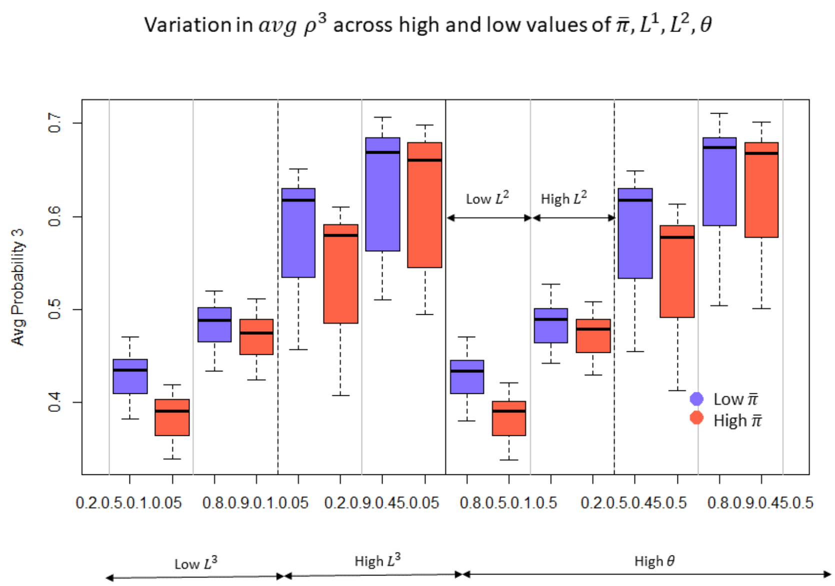

Figure 8a. Variation in percentage of households fully covered

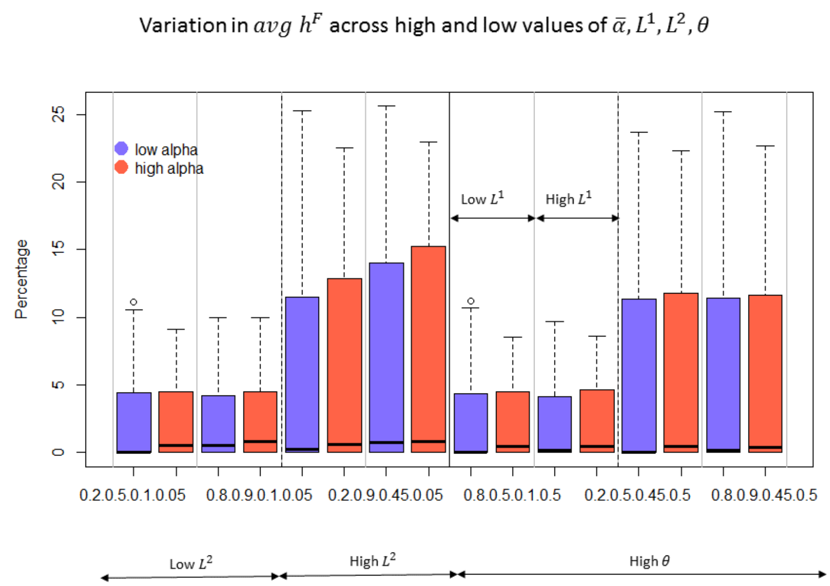

Figure 8b. Variation in percentage of households with binding budget constraints and coverage limits

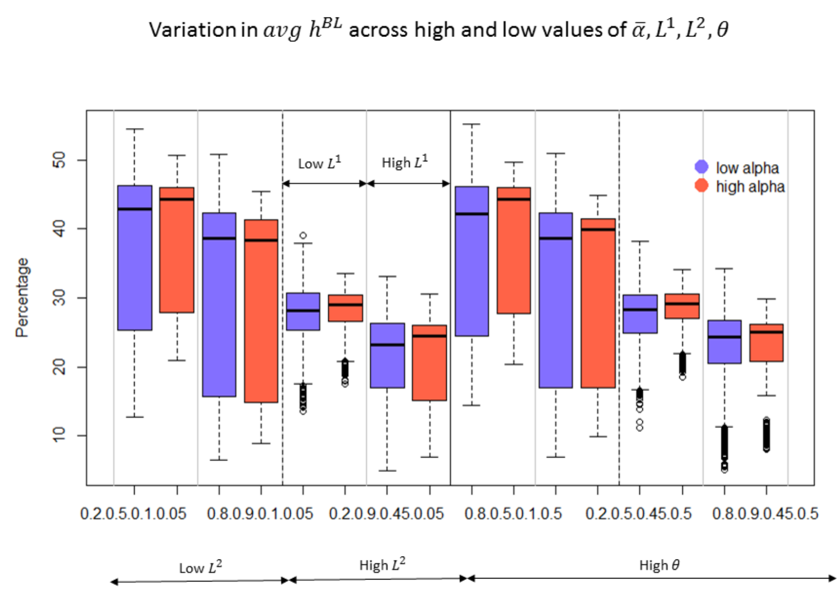

Discussion

This paper tests the viability of creating an energy insurance market for energy self-sufficient households with contract prices according to their preferences for risk, allowing utilities to continue generating revenue from households that no longer rely on electricity bought from the grid. The utility serves the purpose of an insurance provider that supplies a menu of contracts, and the households select the appropriate product based on expected loss, risk profile and budget. In essence, this business model is a long-term contract between the utility and self-sufficient households that helps the utility to reduce the volatility of its revenues, given that their profit margins are unlikely to increase in the future. This is a reflection of new technologies lowering barriers to entry and therefore increasing the possibility of consumers obtaining electricity from other sources.

The static model illustrates the possible nature of new insurance products, the profile of households that purchase them, and the pricing policies set by the utility who have perfect information about the households and operate in one period only. Being uninsured can be a rational choice, even for people who are risk-averse, and is not just an outcome of unaffordability. Individuals may choose to eschew market insurance because of its price relative to that of other goods and services, subjective assessments of personal risk and risk tolerance (Ehrlich and Yin 2017).

The dynamic model, implemented as an agent-based model (ABM), extends the static model to incorporate imperfect information on the part of the utility, as well as the bounded rational nature of households (who may choose to over or under insure), and the role of dynamic learning mechanisms in driving the long-term trajectory of the market. We find that under the general relaxed assumptions of the dynamic model, a stable market can exist, where prices converge to a long-run equilibrium, and the distribution of choices made by households becomes stationary (within limits).

Expected utility theory predicts that people buy insurance that features large deductible and very deep coverage. Myopic loss averters, however, may find it reasonable to pay to mitigate isolated risks for a proportional exorbitant price inconsistent with expected utility theory. Hence in an economy of myopic loss averters, there would be a number of small-scale insurance contracts sold at high premiums, as well as many households left uninsured (Rabin and Thaler 2001). The use of an ABM model allows us to test this anomaly. We find that the coverage limits seem to most influence the choice of contracts at the household level and the final market prices. This would imply that there could be a sizeable market for this service, concurring with the finding that the revenue of the utility-insurer depends mostly on the coverage limits of the contracts and households' potential energy needs. Thus, the utility has a degree of control over its revenues by choosing the coverage limits or, in other words, the riskiness of the contracts available, as well as the nature of the price adjustments it carries out in the market. At the same time, the long-term choices made by households may be influenced by the riskiness associated with the contracts as well as the possible losses they may face.

This proposal is in line with the general trend in regulatory governance that moves towards more flexible, incentive-based and indirect regulation (Parmet 2013). Self-regulation and decentralized mechanisms would be a natural fit in the distributed power generation based sector. We showed that the insurance market proposed in this paper would help to internalize the risk to a great extent, reducing blackouts significantly. However, there may be households unable to cover their full energy needs due to budgetary considerations or the extreme nature of their risk profiles. Our simulations show that, on average, between 1 to 15 percent of those households that would otherwise experience a complete loss of power can fully cover their excess energy needs through insurance. From those households that would otherwise experience a complete loss of power, between 50 to 70 percent are budget constrained and would still be able to partially cover their excess energy needs. Further, the convergence results imply that, at least in theory, such markets can be stable in the medium- to long-run. This means that regulators and utilities could include this as a long-term scenario when exploring business models should distributed power systems become more widely used. This paper shows a feasible way forward for both utilities and regulators in the event of widespread distributed power systems.

The model throws up several interesting points and directions for future research. The current model does not restrict the possible prices, nor does it restrict the nature of updates of prices and coverage limits. Such constraints may change the evolutionary trajectory of the market significantly, especially given the interactions we saw between the relevant parameters.

The second crucial point for regulation would be in the nature of competition in the model. The current model only allows for one utility. However, we can easily envisage scenarios where utilities compete in the market to provide coverage. Allowing for competition between utilities will certainly generate interesting insights on how market shares evolve under various conditions.

Conclusion

In central electricity systems, planners invest in capacity higher than the peak load in order to have reserve margins that can deliver almost perfect coverage, as they can spread this cost across all customers. This approach was more feasible when technology options were limited, and the underlying assumption was that blackouts have an infinite negative value.

In a decentralized and distributed power system, however, this approach seems unnecessarily expensive and unfair, as all customers pay in equal terms regardless of their risk preferences. It also has the potential to be self-defeating for utilities as they would no longer be able to spread fixed reliability costs across a large base of customers if a growing share of them decide to leave the grid.

To deal with security of supply in this context, this paper tested the creation of a reliability insurance market where households can decide their level of protection according to their preferences and pay accordingly. We find it is more efficient for households to transfer the 'last mile' of risk to the utility rather than bear the disutility of a blackout. We find that an insurance market can act as an indirect regulatory mechanism to manage reliability in a distributed power market.

Acknowledgments

The authors gratefully acknowledge the comments and contributions from participants of the workshop "Integrated Approaches to Decentralized Electricity Transitions," hosted by KAPSARC on October 4, 2018, in Brussels, Belgium. The authors are grateful to Iqbal Adjali, Amro Elshurafa, Jorge Blazquez, Baltasar Manzano and Fatih Karanfil for comments on earlier versions of this paper. All remaining errors are the authors'.

References

Berkes, Fikret. 2010. "Devolution of environment and resources governance: trends and future." *Environmental Conservation* 37, 4: 489-500

Calhoun, Koben, Iain Campbell and Doug Miller. 2017. "Lumens as a Service: How to Capture the Technology-Enabled Business Opportunity for Advanced Lighting in Commercial Buildings." Rocky Mountain Institute, May 4.

Ehrlich, Isaac, and Yong Yin. 2018. "The problem of the uninsured." *Research in Economics* 72, 1: 147-168.

Felder, Frank A., and Rasika Athawale. 2014. "The life and death of the utility death spiral." *The Electricity Journal* 27, 6: 9-16.

Fuentes-Bracamontes, Rolando. 2016. "Is unbundling electricity services the way forward for the power sector?." *The Electricity Journal* 29, 9: 16-20.

Fuentes, Rolando, Jorge Blazquez and Iqbal Adjali, I. 2019. From vertical to horizontal unbundling: A downstream electricity reliability insurance business model. *Energy Policy* (forthcoming, accepted for publication Febraury 26th, 2019)

Gosden, Emily. 2016. "Households could be charged annual 'insurance premium' for access to electricity grid." *The Telegraph*. May 29.

Green, Richard, and Iain Staffell. 2017. "Prosumage" and the British electricity market." *Economics of Energy & Environmental Policy* 6, 1: 33-49.

Muaafa, Mohammed, Iqbal Adjali, Patrick Bean, Rolando Fuentes, Steven O. Kimbrough, and Frederic H. Murphy.2017. "Can adoption of rooftop solar panels trigger a utility death spiral? A tale of two US cities." *Energy Research & Social Science* 34: 154-162.

Parmet, Wendy E. 2013. "Beyond paternalism: rethinking the limits of public health law." *Conn. L. Rev*. 46: 1771.

Pearce, David. 2002. "The insurance industry and the conservation of biological diversity: An analysis of the prospects for market creation". Organisation for Economic Co-operation and Development (OECD): Paris,

Rabin, Matthew, and Richard H. Thaler. 2001. "Anomalies: risk aversion." *Journal of Economic perspectives* 15, 1: 219-232.

Watkins, Jonathan. 2017. "Five promising consumer business models to transform low carbon heating and well-being at home." Catapult Energy Systems

Technical appendix

Static model

The household

We first examine a representative household's choice of insurance coverage from an energy supplier. We consider a simple one-period model of a representative electricity market, with N households and a single supplier S. Each household $n \in N$ has installed energy generation capacity which provides a fixed $\bar{C}_n$ units of energy at any given period. A representative household's energy demand C_n is assumed to be stochastic, such that $C_n \leq \bar{C}_n$ with probability $1 - \pi$ and $C_n > \bar{C}_n$ with probability π.

Each household is characterized by a 'dis-utility' function $U(L_n)$ where L_n is the energy loss or shortfall faced by a household n, defined as:

$$L_n = \begin{cases} 0, & if\ C_n \leq \bar{C}_n \\ C_n - \bar{C}_n, & otherwise \end{cases} \quad (1)$$

Households are considered to be risk-averse. We define the (dis-)utility of loss L_n for household n as a monotonically decreasing strictly concave function:

$$U(L_n; \alpha_n) = \bar{U} - e^{\alpha_n L_n} \quad (2)$$

$$where,\ \bar{U} > 1\ and\ \alpha_n > 0.$$

Note that the first and second derivatives of the utility function specified above are $U'(L_n) = -\alpha e^{\alpha L_n} < 0$ and $U''(L_n) = -\alpha^2 e^{\alpha L_n} < 0$, respectively. The above specification implies that the household specific coefficient of absolute risk aversion (CARA) in our model is given by:

$$\frac{U''(L_n)}{U'(L_n)} = \alpha_n \quad (3)$$

Each household produces energy at full capacity at every period, but its energy demand may exceed installed capacity with probability π as specified above. If the demand exceeds installed capacity, the household has the option of purchasing energy from the supplier at a contracted two-part tariff (p, F), where p is the per unit price of energy bought, and F is the lump-sum 'standing charge.' Each household is endowed with a specific budget B_n dedicated for additional energy needs (including insurance coverage), and hence faces a budget constraint:

$$p(L_n - L_n^*) \leq B - F \tag{4}$$

where $(L_n - L_n^*)$ is the 'coverage' provided by the supplier to household n in case of a shortfall in its domestic production. Consequently, L_n^* is the level of shortfall or loss in energy that the household is willing to face, given its budget constraint and utility specification. In other words, L_n^* is the household specific loss not covered by the insurance contract with the supplier. This specification allows the household to draw as much energy as it wants from the supplier, provided it pays the requisite price[5].

Unlike the standard microeconomic choice specification of 'energy' versus 'all other goods,' we adopt a more straightforward approach where the household has already decided to allocate a budget. This makes the model more tractable without any loss of generality, as the central question is on how insurance contracts should be structured rather than on deriving a household's energy demand. Also, it is increasingly being recognized that energy demand exists not for its own sake, but as it acts as a medium to consume other goods and services (heating/cooling, lighting and so forth)[6]. Thus, modeling energy (derived) demand needs to explicitly consider the trade-off between the consumption of goods and services using this energy against all others. However, using such a specification needlessly complicates the model, given that the demand for insurance is itself a second order derived demand based on the demand for energy. Thus henceforth, household-specific budgets for insurance or additional energy needs are assumed to be exogenously fixed at B.

Given that households are symmetric, we drop the subscript n from the analysis which follows.

For the sake of simplicity, assume that households are aware of the exact level of loss that they can incur in case installed capacity falls short of demand and there is no insurance coverage. This implies, *ex ante*, households have full knowledge of the level of loss that they might incur (say, $\overline{L} > 0$) and the probability of such an outcome (π), but not whether the outcome will actually take

[5] We relax this assumption with a coverage limit in the dynamic model, in the next section. The presence or absence of this additional constraint has no influence on main results in the static framework.

[6] See Hunt and Ryan (2015), Walker and Wirl (1993), Goerlich and Wirl (2012), who argue among similar lines.

place. Note that this is qualitatively similar to the general situation where any positive level of loss can be incurred with a known continuous probability over the support (-$\bar{C}, \overline{L}$).

Figure 1. Exponentially decreasing utility function and the certainty equivalent.

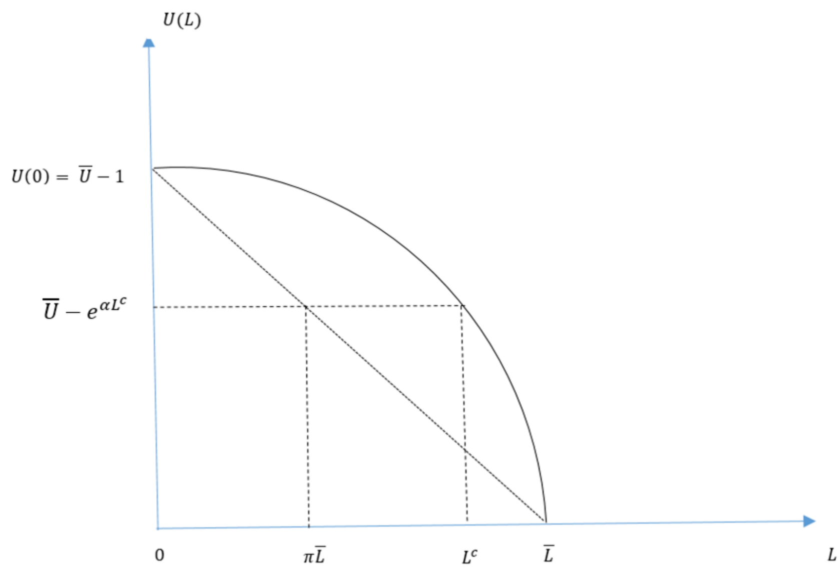

We are now ready to state the main results of this section pertaining to the equilibrium decisions of a representative household and the supplier.

Lemma 1: A rational household characterized by risk aversion α and potential loss $\overline{L}$ will be willing to accept an insurance contract offering a cover of κ only if

$$\kappa \geq \overline{L} - \frac{1}{\alpha}\log[(1-\pi) + \pi e^{\alpha\overline{L}}]$$

Proof: Let L^c be the 'certainty equivalent' amount of loss for the household. Hence by definition,

$$U(L^c) = (1-\pi)U(0) + \pi U(\overline{L})$$

$$\text{or,} \quad \overline{U} - e^{\alpha L^c} = (1-\pi)\big(\overline{U} - 1\big) + \pi(\overline{U} - e^{\alpha\overline{L}})$$

Solving the above for L^c, we get

$$L^c = \frac{1}{\alpha}\log\left[(1-\pi) + \pi e^{\alpha\bar{L}}\right], \quad (5)$$

which is the level of 'certain' loss that makes a risk-averse household indifferent to either accepting or rejecting the random loss (see Figure 1).

If an insurance contract provides a cover κ, the household will either face a loss of 0 (with probability $1-\pi$) or a loss of $\bar{L}-\kappa$ (with probability π). If $\bar{L}-\kappa > L^c$, then for a monotonic utility function $U(\bar{L}-\kappa) < U(L^c)$. In other words, the expected utility from the contract is less than the expected utility of a household that does not accept the contract and takes on the full risk. Hence, for a rational household to accept an insurance contract, $\bar{L}-\kappa \leq L^c = \frac{1}{\alpha}\log\left[(1-\pi) + \pi e^{\alpha\bar{L}}\right]$. The proof follows.

Thus Lemma 1 introduces a lower bound on the loss that a household will be willing to accept *in the case where the household has insurance cover*. The next proposition defines the equilibrium level of cover that the household purchases.

Proposition 1: Facing an insurance cost of (p, F), where $B > F$, a rational utility maximizing household will opt for an equilibrium level of coverage κ^*, such that,

$$\kappa^* = \begin{cases} \dfrac{B-F}{p}, & if\ \bar{L}-\kappa^* \leq L^c \\ 0, & otherwise \end{cases}$$

Proof: The single period optimization problem faced by the household is the following:

$$\max_L\ (1-\pi)U(0) + \pi U(L)\ s.t.\ L \leq L^c\ \ and\ p(\bar{L}-L) \leq B-F$$

Since $U(0) = \bar{U}-1$, which is a constant, the problem reduces to maximizing $\pi U(L)$ wrt L subject to the incentive and budget constraints, respectively. Now as $U'(L) < 0$ for all $L \geq 0$, and as the budget B is specific to energy consumption only, we must have $\left(\bar{L} - L^*(p,F)\right) = \frac{B-F}{p}$, as long as the incentive constraint $L^*(p,F) \leq L^c$ holds (Lemma 1). Replacing $\kappa^* = \bar{L} - L^*(p,F)$ above, proof of the first part follows. Now suppose the incentive constraint does not hold, that is, $\bar{L}-\kappa^* > L^c$. In this situation, the household would be better off facing the risk without any insurance cover (Lemma 1), implying $\kappa^* = 0$. The proof follows.

Proposition 1 implies that the household uses up the entire allocated budget to purchase the optimum level of cover. However, as a corollary to Proposition 1, this implies that the price of insurance (p, F) should not be too high, causing κ^* to be so low that $\overline{L} - \kappa^* > L^c$, in which case the household would prefer not to buy the insurance contract.

The supplier

The results above explore the choices made by households with respect to the level of coverage they buy from the supplier. The households themselves are characterized by the energy budget (B), risk aversion (α), the probability of facing a loss (π), and the shortfall in energy supply they can potentially face $(\overline{L})$. In this section, we identify what the 'ideal' contract is for a household, from the perspective of the supplier.

Suppose that the supplier faces a constant marginal cost c and a fixed cost Φ of including a household within the grid and supplying it with energy. The following proposition characterises the optimum pricing strategy for the supplier for a specific household where the supplier posseses full and complete information about the household.

Proposition 2: Suppose that the supplier possesses full and complete information about a household's energy budget (B), risk aversion (α), the probability of facing a loss (π), and the potential shortfall in supply $(\overline{L})$. In equilibrium, the supplier can maximize its profit from this household by offering it a contract (p^*, F^*), where

$$p^* = \frac{B - F^*}{\overline{L} - L^c},$$

$$F^* = \begin{cases} \phi, & if\ \phi < B \\ 0, & otherwise \end{cases}$$

$$L^c = \frac{1}{\alpha} log\left[(1 - \pi) + \pi e^{\alpha \overline{L}}\right],$$

as long as, $p^* > c$. A rational utility maximizing household will accept such a contract as it satisfies both the budget and incentive constraints.

Proof: For any given contract (p, F) being offered by the supplier, the expected demand from a household for the supplier's energy is:

$$\frac{\pi(B-F)}{p}$$

If the contract is accepted, the expected profit Π_s for the supplier from this household, given marginal cost c and fixed cost Φ, is:

$$\Pi_s = (p-c)\pi\frac{(B-F)}{p} - \Phi + \mathrm{F}$$

The supplier's optimization problem is to $\max_{\mathrm{p,F}} \Pi_s$. Now,

$$\frac{\partial \Pi_u}{\partial p} = \frac{c\pi(B-F)}{p^2} > 0$$

Hence, the supplier should raise its per unit price p as long as either the budget or the incentive constraint of the household becomes binding.

In equilibrium,

$$L^*(p,F) \leq L^c \qquad \text{(from Lemma 1),}$$

$$\text{and } \left(\overline{L} - L^*(p,F)\right) = \frac{B-F}{p} \qquad \text{(from Proposition 1),}$$

where, $L^c = \frac{1}{\alpha}\log\left[(1-\pi) + \pi e^{\alpha\overline{L}}\right]$.

Substituting $L^* = \overline{L} - \frac{B-F}{p}$ from the second expression into the first, we have

$$p \leq \frac{B-F}{\overline{L} - L^c}$$

As $\frac{\partial \Pi_u}{\partial p} > 0$, the profit maximizing price given any F is,

$$p^* = \frac{B-F}{\overline{L} - L^c}$$

Given that $\frac{\partial \Pi_u}{\partial F} = -\frac{p-c}{p} < 0$ if $p > c$, F can then be benchmarked to the fixed cost ϕ of supplying energy to the household with budget B. If the household has a sufficient budget to cover the fixed costs, set $F^* = \phi$, otherwise set $F^* = 0$. In both cases, the equilibrium profit $\Pi^* = B - \phi -$

$c(\overline{L} - L^c)$, and hence the supplier is indifferent between a positive or a zero lump sum component. This completes the proof.

Figure 2 illustrates the optimum price setting principle. The optimum unit price is higher for households that do not pay the lump sum ($F^* = 0$), and lower for households that pay a positive lump sum ($F^* = \phi$).

Note that the above analysis assumes that the supplier has full knowledge of a household's characteristics such as its budget, levels of risk aversion and the probability of loss. In the more realistic case where it does not and is only aware of the probability distribution of each of these characteristics in the market, the above result on optimum pricing strategy may be modified to reflect the expected values of these characteristics in the market. In such a case, the supplier cannot fully discriminate between households but may offer a single contract which only eligible households (for whom the constraints are satisfied) will choose. Alternatively, it can offer a menu of contracts, each with a different level of per unit price and lump sum, and households will select the contract which maximizes their utility. In both cases, it is possible that some households do not purchase any contract, thus choosing to stay 'dark' when there is a shortfall in their self-generated energy supply.

Figure 2. Plot of Π_u as a function of p, given the budget and incentive constraints of a household. The optimum p^* is also indicated under the condition $p^* > c$.

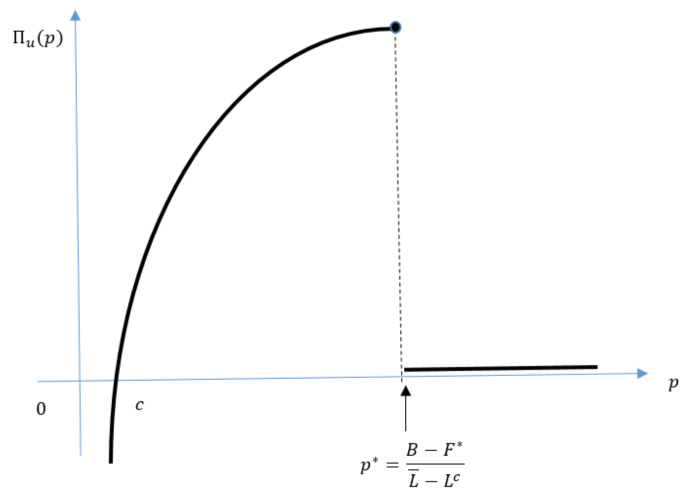

Dynamic Model

We assume that time extends from 0 to T, where T is sufficiently large so that the supplier and the households are not able to factor it explicitly into their optimization. We also assume that households have complete information about the past and no information about the future. Their energy demand for period t is uniquely determined in t and is unconditional on any previous period's usage (no seasonality or autocorrelation). The supplier also only provides energy coverage for a single period. All assumptions regarding household behaviour hold true in this model for any specific period $t \in T$. To start, we only consider rational households and suppliers with no learning or adaptation capabilities.

Given that the period-specific conditions for a household remain the same as in the static framework, and the insurance coverage provided by the utility is essentially for a single period (and cannot spill over into the future), the household's rational equilibrium demand $\kappa^*(t)$ in period t is identical to that characterised in Proposition 1.

$$\kappa_t^* = \begin{cases} \dfrac{B - F_t}{p_t}, & if \;\; \bar{L} - \kappa_t^* \leq L^c \\ 0, & otherwise \end{cases}$$

The supplier's problem is to maximize the discounted sum of period profits $\Pi_t(p_t, F_t)$ over $t \in T$, where $T \to \infty$. Given that there are no inter-temporal constraints or spill-overs in the model, the inter-temporal maximization is equivalent to the period maximization problem, implying that the supplier maximizes per period profit in the dynamic model as well. Hence, in this version of the dynamic model, the equilibrium p_t^* and F_t^* are identical to the static equilibrium values expressed in Proposition 2.

Adaptation and learning

We now introduce an alternative dynamic specification, where both the supplier and the households are allowed to learn from past behavior and outcomes and adapt their behaviour over time. This is specifically done in order to reduce the rationality burden on both, as is required by the economic model specified earlier. This allows households to make 'mistakes' in their choice of insurance contract (in terms of coverage purchased) and the ability of the utility to change the tariff charged in each contract offered based on demand. This makes the modeling less restrictive and provides us with a more flexible implementation of the agent-based model.

In this specification, neither the supplier nor the households are assumed to be rational, but alternative characterizations of behavior are benchmarked against the rational equilibrium outcome. We make the following adjustments to the model to incorporate learning behavior.

First, we introduce a coverage limit in the contract, which acts as an additional constraint for households in the model. This implies that an energy insurance contract is now of the form (p, F, L), where p and F represent the per unit and lump sum prices respectively, as before. The additional L represents the maximum coverage allowed within the contract at the given prices.

Second, we allow for the possibility of multiple insurance contracts being made available by the supplier to the households. In this scenario, the supplier provides a menu of alternative K contracts $\{(p_t^1, F^1, L^1), \dots, (p_t^K, F^K, L^K)\}$ from which a household has to select one.

Third, we introduce an adaptation rate parameter τ, which represents how quickly households and the supplier are allowed to change their demand and supply behaviours, respectively. At time $t = 0$, the supplier offers the contracts $\{(p_0^1, F^1, L^1), \dots, (p_0^K, F^K, L^K)\}$, from which each household selects one. We assume that the supplier updates the contracts every τ periods, and the households, once they have chosen a contract, are contractually bound to it for τ periods as well. After every τ

periods, the supplier has the option to modify its offerings and the households have the option to update their choices. These sets of contracts and choices once again remain fixed for the next τ periods. Thus the parameter τ indicates the rate at which actors in the market are able to update their decisions, with a smaller value indicating a faster update[7]. We now define the updated rules of this framework that households and the supplier follow.

The households choose from the menu probabilistically. Hence, household n is characterised by a probability distribution $(\rho^1(t), \dots, \rho^k(t))_n$ in period t, such that in every τ periods it selects contract $k \in \{1, \dots, K\}$ with probability ρ_t^k. While a household chooses a contract once every τ periods, the probability distribution is updated in every period t. Hence, at the time of choosing the next contract, the household faces a potentially new probability distribution over the choices representing the cumulative effect of the last τ periods. The dynamic update of a household's probability distribution over the available menu of contracts at period t is assumed to follow a simple reinforcement learning algorithm based on which of the three constraints had been binding in period $t - 1$. These are summarized as follows:

(Rule 1H) If the budget constraint was binding in period t, in period $t + 1$ reduce the probability weights on contracts with relatively higher per unit prices and increase the probability weights on contracts with relatively lower per unit prices.

(Rule 2H) If the incentive constraint was binding in period t, in period $t + 1$ reduce the probability weights on contracts with relatively lower coverage limits and increase probability weights on contracts with relatively higher coverage limits.

(Rule 3H) If the coverage limit constraint was binding in period t, in period $t + 1$ reduce the probability weights on contracts with relatively lower coverage limits and increase probability weights on contracts with the highest coverage limits.

Note that while Rules 2 and 3 are qualitatively similar, applicable under the general condition of insufficient coverage in the previously chosen contracts, Rule 3 encourages a faster movement

[7] As an example, one can consider each period to be a month and if $\tau = 12$, each contract lasts for 12 months. Households choose a new contract annually. Note that the annual nature of contracts is just an example, and the model update can happen faster or slower as desired.

than Rule 2 towards contracts with the highest coverage limits. Exactly how Rules 1-3 are operationalized depends on the number of contracts K, and will be discussed in the next subsection.

The supplier is allowed to adjust the *unit prices offered* in the contracts every τ periods, reflecting changing demand conditions. Let $D_k^*(t)$ be the total number of households subscribing to contract k in period t, and let $0 < \theta < 1$ be a threshold parameter defined exogenously. The update rule for any contract $k \in \{1, \dots, K\}$ can then be stated as:

(Rule 1S) Every τ periods, change unit price of contract k from p^k to $p^k(1 + \Delta p^k)$, where Δp^k is defined as:

$$\Delta p^k = \begin{cases} \dfrac{D_k^*(t) - D_k^*(t-\tau)}{D_k^*(t-\tau)} & if, \quad \left|\dfrac{D_k^*(t) - D_k^*(t-\tau)}{D_k^*(t-\tau)}\right| > \theta \\ 0, & otherwise \end{cases}.$$

The above rule makes price adjustments in the dynamic model demand driven, where existing per unit price is increased (or decreased) in the same proportion as the change in cumulative demand over τ periods, provided the change is large enough, as determined by the parameter θ. Note that a higher θ implies a *less frequent* price update than a lower θ, but each update is by a greater amount. On the other hand, a smaller θ implies more frequent updates than a higher θ but by a smaller proportion. This rule internalises the update amount and makes it completely demand driven.

Agent-based framework

Each household agent has the following parameters: $\alpha, \bar{L}, B, \pi$, randomly drawn from uniform distributions: $\alpha \in (0, \bar{\alpha}), \bar{L} \in (0, \bar{\bar{L}}), B \in (0, \bar{B}), \pi \in (0, \bar{\pi})$. The upper limits of the supports are parameters in the simulations, details of which are provided in Table 1. The supplier provides three alternative contracts for households to select from, i.e., $K = 3$. The households choose a contract and the menu of contracts is updated every 12 steps within the simulation ($\tau = 12$). A household at $t = 0$ starts with an even probability distribution (0.33, 0.33, 0.33) across all contracts in the menu.

By default, we set contract 1 to be a 'spot' contract with no pre-specified energy limit, i.e., potentially $L^1 \to \infty$. Contracts 2 and 3 set upper limits on how much energy a household can draw. We impose the following constraints on the contracts: $L^1 > L^2 > L^3$ and $p^1 > p^2 > p^3$. Thus the spot contract is the risk-less contract but with the highest per unit price, while contract 3 is the riskiest contract but with the lowest per unit price. Contract 2 is an intermediate one, providing a balance between risk and cost.

As described above, the choice probabilities are updated by the household at the end of period t based on which of the three constraints are binding on the household in the same period. For any given household, let the actual contract chosen in period t be denoted by $k^*(t)$. Let δ be an adjustment parameter, which is the amount households adjust their choice probabilities every period. In such a case, Rules 1H, 2H, 3H are operationalized in the following manner:

Rule 1H - Budget constraint binds in period $t-1$

a. If $k^*(t-1) = 1 \rightarrow \left\{ \begin{array}{c} \rho^1(t) = \rho^1(t-1) - \delta, \\ \rho^2(t) = \rho^2(t-1) + \frac{\delta}{2} \\ \rho^3(t) = \rho^3(t-1) + \frac{\delta}{2} \end{array} \right\}$

b. If $k^*(t-1) = 2 \rightarrow \left\{ \begin{array}{c} \rho^1(t) = \rho^1(t-1) - \frac{\delta}{2}, \\ \rho^2(t) = \rho^2(t-1) - \frac{\delta}{2} \\ \rho^3(t) = \rho^3(t-1) + \delta \end{array} \right\}$

c. If $k^*(t-1) = 3 \rightarrow \left\{ \begin{array}{c} \rho^1(t) = \rho^1(t-1) - \frac{\delta}{2}, \\ \rho^2(t) = \rho^2(t-1) - \frac{\delta}{2} \\ \rho^3(t) = \rho^3(t-1) + \delta \end{array} \right\}$

Rule 2H – Incentive constraint binds in period $t-1$

a. If $k^*(t-1) = 1 \rightarrow \left\{ \begin{array}{c} \rho^1(t) = \rho^1(t-1) + \delta, \\ \rho^2(t) = \rho^2(t-1) - \frac{\delta}{2} \\ \rho^3(t) = \rho^3(t-1) - \frac{\delta}{2} \end{array} \right\}$

b. If $k^*(t-1) = 2 \rightarrow \left\{ \begin{array}{c} \rho^1(t) = \rho^1(t-1) + \frac{\delta}{2}, \\ \rho^2(t) = \rho^2(t-1) + \frac{\delta}{2} \\ \rho^3(t) = \rho^3(t-1) - \delta \end{array} \right\}$

c. If $k^*(t-1) = 3 \rightarrow \left\{ \begin{array}{c} \rho^1(t) = \rho^1(t-1) \\ \rho^2(t) = \rho^2(t-1) + \delta \\ \rho^3(t) = \rho^3(t-1) - \delta \end{array} \right\}$

Rule 3H – Coverage limit constraint binds in period $t-1$

a. If $k^*(t-1) = 1 \rightarrow \begin{Bmatrix} \rho^1(t) = \rho^1(t-1) + \delta, \\ \rho^2(t) = \rho^2(t-1) - \frac{\delta}{2} \\ \rho^3(t) = \rho^3(t-1) - \frac{\delta}{2} \end{Bmatrix}$

b. If $k^*(t-1) = 2 \rightarrow \begin{Bmatrix} \rho^1(t) = \rho^1(t-1) + \delta, \\ \rho^2(t) = \rho^2(t-1) - \frac{\delta}{2} \\ \rho^3(t) = \rho^3(t-1) - \frac{\delta}{2} \end{Bmatrix}$

c. If $k^*(t-1) = 3 \rightarrow \begin{Bmatrix} \rho^1(t) = \rho^1(t-1) + \frac{\delta}{2} \\ \rho^2(t) = \rho^2(t-1) + \frac{\delta}{2} \\ \rho^3(t) = \rho^3(t-1) - \delta \end{Bmatrix}$

Both Rules 2H and 3H are similar in that they both add weight on contracts with higher coverage. The only distinction between them is in the rate at which they shift the households' choice towards the spot contract: Rule 3H shifts the probability weights in favor of the spot contract at a higher rate than 2H, while Rule 1H shifts the weights in favor of cheaper contracts. The adjustment parameter δ is fixed exogenously at 0.05. Once the updates have taken place, the probabilities are rebased to ensure that they are bounded within the [0, 1] interval.