An empirical model of the decision to switch between electricity price contracts*

Gauthier Lanot
Department of Economics
Umeå School of Business and Economics,
Umeå University, Sweden, gauthier.lanot@umu.se,

Mattias Vesterberg
Department of Economics
Umeå School of Business and Economics,
Umeå University, Sweden, mattias.vesterberg@umu.se,

We present a novel model for a time series of individual binary decisions which depends on the history of prices. The model is based on the Bayesian learning procedure which is at the core of sequential decision making.

We show that the model capture dependence on past events and past priors in a straightforward fashion, the model capture some dependence on initial condition, here in the form of the prior at the start of the decision period, and that estimation through maximum likelihood is straightforward.

We estimate the parameters of the model on a sample of Swedish households who have to decide over time between competing electricity contracts. The estimated parameters suggest that households respond to prices by switching between contracts, and that the response can be rather substantial for alternative price processes.

Key words: Price, Contract Choice, Bayesian Learning, Time Series, Binary Decision, Survival analysis

1. Introduction

Households repeatedly face the decision to switch between many different types of contracts and tariffs, including electricity contracts (Goett et al. (2000); Juliusson et al. (2007)), internet services (Lambrecht et al. (2007)), telephone services (Miravete (2002); Braun and Schweidel (2011)) and mortgage rates (Dhillon et al. (1987); Campbell and Cocco (2003)). Typically, these choices involve deciding whether to remain on the current contract, or to switch to another, possibly more beneficial, contract. Analysing these choices is important from an economic perspective as it provides insights into the precise nature of the household's decision making, how they process the information available to them and how their decision depends on their current situation. The financial implications of the choices are important for individual customers and retailers (Lambrecht and

* The authors would like to thank David Granlund, Bertrand Koebel and seminar participants at The Economic Policy Network for helpful comments and suggestions

Skiera (2006); Gupta et al. (2006)). Furthermore, the optimal regulation of such markets rests on the results from analyses of this kind (Brennan (2007)) and this is true in particular for the case of electricity contracts, which is the topic of the current paper.

Since the deregulation of the Swedish electricity retail market in 1996, households are free to choose between hundreds of retailers, each offering several different contracts. Roughly 40 percent of all Swedish households have chosen a contract with prices per kWh fixed for a year or longer (fixed-price contract), and another 40 percent have chosen a contract where prices vary by month (variable-price contract). Most of the remaining households are on so-called default contracts with the incumbent retailer, which households automatically are assigned to if they do not make an active choice to either a variable-price or fixed-price contract. These contracts typically have prices varying by season, and the price per kWh is higher than both variable-price and fixed-price contracts.

The choice of electricity contract is one of the many margins of choice the household face (together with choice of retailer, energy efficiency, etc.), and it is the aggregate of these decisions that determines the overall demand for electricity. For example, a household choosing a fixed-price contract will obviously not respond to prices in the short run (i.e., prices are fixed within a contract period). Failing to consider this when estimating the aggregate demand for electricity could potentially underestimate the demand elasticity, and households may be more price sensitive than previously thought once we correctly condition on contract choice. Furthermore, on a competitive market as the Swedish retail market, cost savings are expected to be larger from switching between contract types than from switching between retailers because prices are similar across retailers but differs between contract types (Littlechild (2006)). This suggests that the choice of contract may have the largest financial implications for electricity consumers, and that the contract switching response is a non-negligible aspect of electricity demand price responsiveness. Understanding the choice between electricity contracts is also of importance to retailers, because this will affect the composition of risk they bear: for fixed-price contracts, it is the retailer that carries the price risk entirely, whereas the household carries the risk for variable-price contracts.

Understanding this choice involves many complicating factors. First, it seems likely that households respond to a sequence of prices, rather than a single price. Second, from the households' point of view, determining which type of contract is the most beneficial to households is a difficult endeavour since future relative prices are unknown and previous relative prices may for some periods be similar and for other periods very different (see Figure 1). Third, although various websites provide price information and assistance for switching between electricity contracts, there may still exist monetary and non-monetary transaction costs to switching (Klemperer (1987); Juliusson et al. (2007); Ek and Söderholm (2008); Ericson (2011)). Hence, if it is costly to switch between contracts, households will avoid frequent switching. The households then need to predict prices in the far future, further complicating the (understanding of the) decision process. As a result, households may abstain from switching between contracts, even if prices ex-post may appear to be in favour of frequent switching. This is true for electricity contracts and most other similar contract choices.

Indeed, consumers have been shown to choose seemingly sub-optimally in many ways, including making choices too infrequently (e.g., electricity retailer choice; Wilson and Price (2010)), and tending to stick with the unfavourable (to them) default choices (Benartzi and Thaler (2007); Fowlie et al. (2015)). The literature on the so-called flat-rate bias, where consumers tend to choose flat-rate contracts when per-use contracts would be better, may be yet another indication of the complexity of these type of choices. For example, Miravete (2002, 2003) studies the theoretical and econometric implications of agents facing uncertainty concerning their future consumption of telephone services, where the uncertainty is resolved through individual and privately known shocks. He finds that consumers switch between tariff options with the explicit goal of minimizing the cost of their service, and they do so in the short term and in response to small differences in billing cost. However, beliefs may be biased, and consumers often choose and retain the wrong tariff. Similarly, Vigna and Malmendier (2006) present evidence that biased beliefs about future behaviour best explain the flat-rate bias for health clubs. This bias also explored in Goettler and Clay (2011), who propose a Bayesian learning model of tariff and usage choice in the online grocer market, and illustrate how their model explains this bias even for rational consumers. Lambrecht et al. (2007) find both flat-rate bias and pay-per-use bias in the context of internet access, and find evidence that the underestimation of usage is a major cause of the pay-per-use bias.

For the case of switching between electricity contracts, descriptive statistics from public data suggest that the response to prices by switching between contracts is not obvious (SCB (2014)). Rather, as is illustrated in Figure 1, even though the price differential between variable-price and fixed-price contracts has varied substantially during the last ten years, few households switch between contracts in the short run. In particular, prices are much more variable than the share of households on each type of contract.[1]

Previous literature on determinants of electricity contract choice is limited, and relies on the stated preference approach. Juliusson et al. (2007) explore how loss aversion and beliefs about future price volatility affect the probability to select a variable-price contract, and find a negative effect for both. Goett et al. (2000) explore the willingness-to-pay for electricity contract attributes, and find that households are willing to pay a small premium (0.008 USD per kWh) to avoid seasonal price variability. There is also some literature on the choice of electricity retailer (see, for example,

[1]Since 2004, the share of households on default contracts has decreased from roughly 60 percent to less than 20 percent.

Figure 1 Number of Swedish households on each type of contract and price per kWh

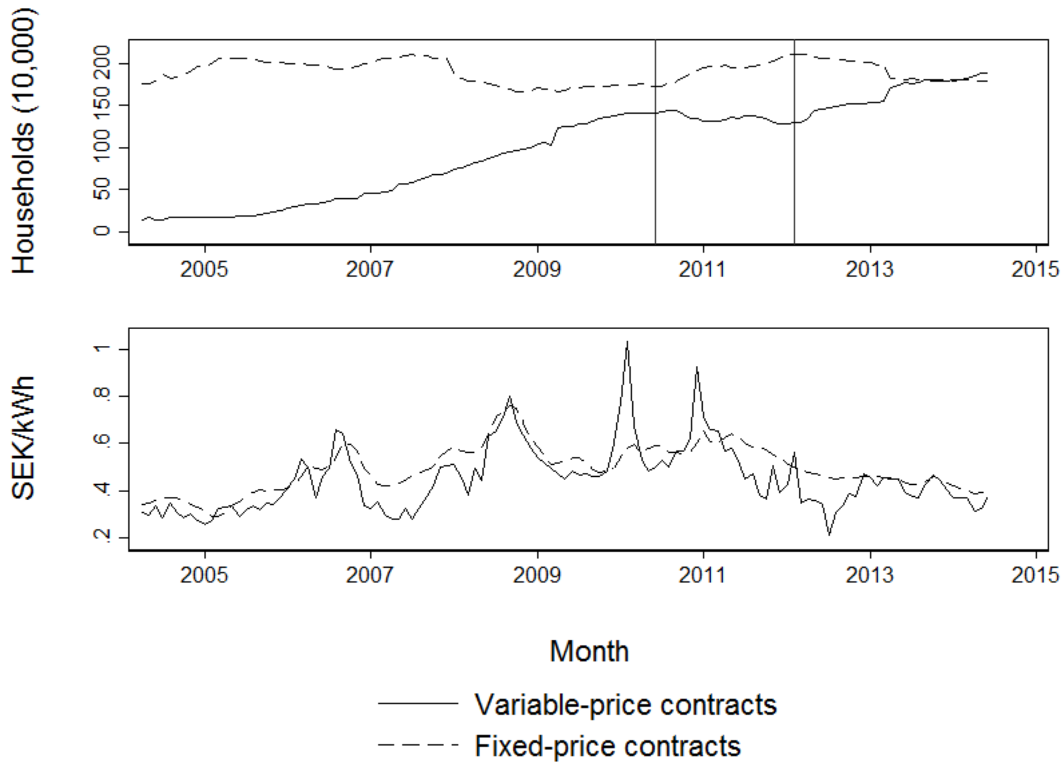

Note: One öre (0.01 SEK) is roughly 0.001 Euro. Fixed-price contract refers to contracts with prices fixed for a year or longer (typically up to three years), and variable-price contract refers to contracts with prices varying by month. In addition to the price per kWh, households pay for transmission, green certificate and taxes. These costs (which are not included in the figure above) amounts to roughly 50 öre/kWh, plus a fixed annual transmission fee varying between 1500 and 15000 SEK depending on the required amperage. These additional costs are identical for both type of contracts. The two vertical lines at June 2010 and February 2012 represents the time period during which households in the sample start their variable-price contract, as explained in Section 3. The households in the sample are then followed over time until June 2014.

Ek and Söderholm (2008) and Ericson (2011)), but it is unclear how the findings in that literature translate into or relate to the decision to switch between electricity contracts.

In this paper, we explore households' decision to switch between electricity contracts in response to current and previous prices. More specifically, we explore the decision to either remain on a variable-price contract or to switch to a fixed-price contract.[2] In particular, we are interested in how sensitive the timing of transition away from the variable-price contract is to current and past prices. To that effect, we propose a model for time series of binary decisions which allow for the

[2]Households on fixed-price contracts face monetary transaction costs if they switch to a variable-price contract before the end of their fixed-price contract term, substantially complicating the analysis. In comparison, households on variable-price contracts are free to switch to a fixed-price contract in any given month.

dependence on the history of prices. Specifically, our model is based on the Bayesian learning procedure which is at the core of sequential decision making. This allows us to capture dependence on a potentially long sequence of information, whereas a simpler logit specification with lagged prices would require the estimation of a large number of parameters to capture this dependence.

The proposed model differs from the much of previous literature on consumer learning (see, for example, Ching et al. (2013) for a literature review). In particular, that literature is often interested in understanding how consumers with imperfect information about contract attributes learn about, e.g., quality, whereas we are interested in a choice where the quality is the same for both contracts but where households are trying to learn which contract that is most preferable to the household.

In short, we find that households do respond to prices by switching between contracts, and that the response varies substantially across price histories experienced by households. Further, we illustrate how the response can be rather substantial for alternative price differential processes. Although the proposed model includes many different parameters and the parameter estimates are not straight-forward to interpret, we illustrate how the results are easily summarized using survival functions.

The contribution of this paper is twofold: first, from an econometric perspective, our model illustrates a novel and relatively straight-forward way of modelling the sensitivity of decision making to new information, conditional on a history of information, and how this can be useful to understand complex decision making. Secondly, as far as we are aware, electricity contract choice is usually not studied using revealed preference approaches where households respond to actual prices. The insights that are generated are important in order to understand household behaviour, and are useful for both retailers and policy makers. While we estimate our model using data on electricity contracts, we argue that both the empirical approach, as well as the findings, are generalizable to other choices.

The rest of this paper is structured as follows: Section 2 describes the empirical model where households update their beliefs about the state of the world as new information becomes available. Section 3 describes the data used to estimate the model. Estimation and results are presented in Section 4, and Section 5 concludes.

2. Empirical model

We wish to model the probability to remain on a variable-price contract at time t against the decision to switch to a fixed price contract given all the information available until then. That is, we propose a model for the conditional probability of the decision $D_t = 1$ to remain with a variable-price contract at time t:

$$\Pr[D_t = 1 | \mathcal{X}_t, t, z_0], \tag{1}$$

given a set of initial household and dwelling characteristics z_0 and a history of information accumulated up to time t, $\mathcal{X}_t$. We interpret this probability as the proportion of households remaining with a variable-price contract at the end of period t among all households with a variable-price contract since $t=1$ who have experienced, for example, a history of price differentials (in the current paper the log of the variable price divided by the fixed price) $\mathcal{X}_t$ since the end of period $t=0$ when their variable-price contract started. In our binary context, the probability to switch to a fixed price contract is simply:

$$\Pr[D_t = 0|\mathcal{X}_t, t, z_0] = 1 - \Pr[D_t = 1|\mathcal{X}_t, t, z_0].$$

For the Swedish retailer market, switching between contracts should not affect electricity consumption levels by any significant amount. First, the only differences between a variable-price contract and a fixed-price contract are the price and the price variability, and relative to the contracts offered in the US, for example, the Swedish contracts do not contain contract-specific attributes such as the increasing block-rate pricing schemes typically found in the US. Second, to a large extent, electricity consumption is determined by temperature and daylight seasonality[3] together with the number and efficiency levels of the appliances used by the household; the former is obviously common to all contracts within broad regions, and the latter is typically fixed in the short run. See, for example, Vesterberg and Kiran B Krishnamurthy (2016) for a detailed description of residential electricity demand in Sweden. Indeed, previous literature on electricity demand typically finds the demand to be inelastic (see, for example, Brännlund et al. (2007) and Krishnamurthy and Kriström (2015) for Sweden). By comparison, the response to outdoor temperature is expected to be substantial, given that heating is by far the largest share of total electricity consumption for most villas in Sweden (Vesterberg and Kiran B Krishnamurthy (2016)). Overall, electricity consumption is believed to be unaffected by the type of contract, and therefore the expenditure variation after a contract switch mostly reflects the price differential. This allows us to model the choice of remaining on a variable-price contract or switch as separate from the choice of how much electricity to consume

Further, we note that the switching costs associated with switching from a variable-price contract to a fixed-price contract should be relatively small. First, households on variable-price contracts are free to switch to a fixed-price contract at any time, and switching in this direction involves no monetary switching costs. In comparison, households on fixed-price contracts have agreed to remain on this contract for the contract period, and typically needs to pay a penalty fee if they wish to cancel their contract in advance. Second, retailers typically offer comprehensive information

[3]Sweden's territory is located in the upper half of the northern hemisphere, between N55° and N70° latitude north. By comparison New York's latitude is N40.7°, Anchorage is N61.2°, while Stockholm and Umeå are respectively N59.3° and N63.8°. Temperatures and daylight hours are therefore very variable over the year and between the north and the south of Sweden.

about prices and how to switch contract on their website. Similarly, the Swedish energy market inspectorate offers similar information on their website (see `www.elpriskollen.se`).

We assume that households understand that the state of the world is binary: in the first case, $S=1$, a variable-price contract is preferable to a fixed-price contract, and $S=0$ if instead the fixed price contract is preferable to the flexible price contract. Households are however are not generally able to identify with certainty the state of the world which describes the data generating process. To form their opinion about the state of the world, households consider the realisation x_t, of a process X_t, which provides some information about the state of the world, S, at time t. In our application, x_t measures the de-seasonalized price differential between the variable-price and the fixed-price electricity unit price at time t.[4] While the current realisation of x_t is informative about the state of the world, knowing x_t given the history of realisations so far is not necessarily enough to identify the state of word exactly. Given $\mathcal{X}_{t-1}$, the observed history of the realisations of the exogenous process X_t since $t=0$ up to time $t-1$, observing x_t is not sufficient to decide whether $S=0$ or $S=1$ with probability 1. We build our empirical model on the following intuition: given the costs and benefits of each type of contract, the household will decide to switch to a fixed-price contract if the evidence against the state of the world, $S=1$, where the variable price is preferable, is strong enough.

In the next sections, we first discuss the issue of signal extraction, i.e. determining the the state of the world given the evidence available . We then propose a model for the probability of a decision to stay or leave a flexible price contract given the evidence available.

2.1. Belief Update

To simplify, we assume that the history $\mathcal{X}_t$ provides households with information about the state of the world (i.e., the sign and magnitude of the expected price differential or its variance) and that, all other things equal, households base their decision in each period on the odds of the state of the world given the history of information so far. We assume that households incorporate new evidence into the odds (the ratio of the probabilities of each state of the world) by following the logic of Bayesian updating, see DeGroot (2005). This relates the posterior odds of $S=1$ against $S=0$ at time t given the history of observations so far, some initial conditions, z_0, and an initial value of the initial odds, o_0, to the product of the prior odds given the observation history until time $t-1$ for the same two events and the likelihood ratio arising from the information revealed at time t by the observation of x_t:

$$\underbrace{\frac{\Pr[S=1|\mathcal{X}_t,t,z_0,o_0]}{\Pr[S=0|\mathcal{X}_t,t,z_0,o_0]}}_{\text{posterior odds},o_t} = \underbrace{\frac{\Pr[S=1|\mathcal{X}_{t-1},t,z_0,o_0]}{\Pr[S=0|\mathcal{X}_{t-1},t,z_0,o_0]}}_{\text{prior odds}} \cdot \underbrace{\frac{\Pr[x_t|S=1,\mathcal{X}_{t-1},t]}{\Pr[x_t|S=0,\mathcal{X}_{t-1},t]}}_{\text{likelihood ratio},\lambda_t}. \tag{2}$$

[4] Because the retail market is competitive and prices are similar across retailers, as explained in the Introduction, there is no need to include the price of competing retailers' contracts

According to this scheme, today's odds of the state of the world given the history of the realisations of the exogenous process, $\mathcal{X}_t$, is expressed as the product of two terms: the first is the odds for the state of the world given the accumulated evidence up to time t, while the second term is the ratio of the conditional probabilities of the current realisation of the exogenous process in either state of the world. The second term therefore captures the relative information contained in x_t about the current state of the world, while the first term measures the relative strength of the household's belief in favour of one or the other state of the world.

If the prior odds at time t are equal to the posterior odds obtained in period $t-1$ after experiencing the information available in $t-1$, x_{t-1}, the relationship between prior and posterior odds allows us to characterise the dynamic of the posterior odds as a function of past evidence. In this case, the dynamic of the log-odds satisfies the relationship:

$$\ln o_t = \ln o_{t-1} + \ln \lambda_t, \tag{3}$$

where o_t stands for the posterior odds in the previous expression, while λ_t stands for the likelihood ratio, i.e. the ratio of the probabilities of the realisation of x_t given either state of the world. The initial log odds, $\ln o_0$, is assumed given. Since the household has chosen a variable-price contract, $\ln o_0$ should be consistent with this choice and we expect $\ln o_0$ to be large and positive.

In this form, the dynamic of the log odds is that of a random walk where $\ln \lambda_t$ plays the role of an innovation and captures the effect of new information contained in x_t on the posterior log odds in t. If $\lambda_t > 1$ then $S = 1$ appears more likely than it was in $t-1$ while if $\lambda_t < 1$ then $S = 0$ appears more likely than it was in $t-1$. Clearly, we can rewrite the log odds at time t as the sum of log likelihood ratios and of the initial log odds:

$$\ln o_t = \sum_{\tau=1}^{t} \ln \lambda_\tau + \ln o_0. \tag{4}$$

An important aspect of the Bayesian information updating scheme is that it provides a natural way to aggregate the information contained in $\mathcal{X}_t$, i.e., provided we know the form of the likelihood ratio, a potentially long sequence of information is summarised completely in the sum of the log-likelihood ratios.

The assumption concerning the relationship between posterior odds at time $t-1$ and prior odds at time t about the state of the world is convenient because of its simplicity. However, because of the random walk structure of the evolution in time of the log-odds, it implies that information acquired some time in the past and information acquired more recently will contribute identically to today's posterior log odds. This may not be a feature of observed information acquiring behaviour. Rather, we may expect recent and distant information to affect today's posterior odds differentially.

We wish to introduce the possibility of some decay in the importance of past information. One way to achieve this is to assume that the prior odds at time t are obtained as a simple transformation of the posterior odds obtained after observing all available information at time $t-1$:

$$\underbrace{\frac{\Pr[S=1|\mathcal{X}_{t-1},t,z_0,o_0]}{\Pr[S=0|\mathcal{X}_{t-1},t,z_0,o_0]}}_{\text{prior odds at time } t} = \exp(\kappa_0)\Bigg(\underbrace{\frac{\Pr[S=1|\mathcal{X}_{t-1},t-1,z_0,o_0]}{\Pr[S=0|\mathcal{X}_{t-1},t-1,z_0,o_0]}}_{\text{posterior odds at time } t-1}\Bigg)^{\kappa_1} \tag{5}$$

where κ_0 and κ_1 are parameters. The transformation of posterior beliefs concerning the state of the world at time $t-1$ into a state of the world at time t is crucial here; without it, we would not be able to characterise the manner in which information accumulates. This assumes that posterior beliefs are modified into prior beliefs by modifying the odds. The LHS of Equation (5) is well behaved, i.e., it takes non-negative values if the input into the power transformation on the RHS can be interpreted as odds.

In particular, Equation (5) modifies the dynamic of the log odds given in (3) to:

$$\ln o_t = \kappa_0 + \kappa_1 \ln o_{t-1} + \ln \lambda_t, \tag{6}$$

and the dynamic of accumulation of information takes the form of autoregressive process of order 1. Hence, a new realisation of x_t adds to the evidence, in the form of $\ln\lambda_t$, to the prior log odds, $\kappa_0 + \kappa_1 \ln o_{t-1}$.

The previous expression describes how, starting from the initial odds, the past information accumulates into today's odds depending on the parameters κ_0 and κ_1[5]:

$$\ln o_t = \ln\lambda_t + \sum_{\tau=1}^{t-1} \kappa_1^{\tau} \ln\lambda_{t-\tau} + \kappa_1^t \ln o_0 + \kappa_0 \frac{1-\kappa_1^t}{1-\kappa_1}. \tag{7}$$

Clearly, the contribution of past information to today's odds decays whenever $|\kappa_1| < 1$. Alternatively, if $|\kappa_1| > 1$, the further in the past an observation the larger its relative weight compared to more recent evidence. Finally, we observe that the dynamic of the log odds introduces a deterministic evolution in time determined by $\frac{1-\kappa_1^t}{1-\kappa_1}$ (which is linear whenever $\kappa = 1$).

2.2. Decision Probabilities

To describe household behaviour, we assume that in the population, the proportion of households that decide to remain with the variable-price contract, given the initial conditions and the evidence accumulated so far, $\Pr[D_t = 1|\mathcal{X}_t, t, z_0, o_0]$, satisfies a logit specification:

$$\Pr[D_t = 1|\mathcal{X}_t, t, z_0, o_0] = \frac{1}{1+\exp\{-z_0\pi_0 - \pi_1 \ln o_t\}} \tag{8}$$

[5]when $\kappa_0 = 0$ and $\kappa_1 \to 1$ this expression agrees with equation (4)

where π_0 and π_1 are parameters which describe the sensitivity of the decision to remain with the variable-price contract on the posterior odds of the state of the world given the information observed at time t. This model is useful since the dependence on the history of the evidence experienced so far is aggregated into $\ln o_t$ as described in equation (7). The aggregation of information depends on the exact expression of the log likelihood ratio as well as on the parameters κ_1 and κ_0. While we expect π_1 to be positive, we have no clear view on its exact magnitude. If both states of the world have equal posterior probabilities, so that $\ln o_t = 0$, the average proportion of exits from variable-price contract is simply $\{1 + \exp(-z_0\pi_0)\}^{-1}$.

This formulation arises naturally in the context of a conventional decision theoretical framework where the expected costs of each decision, to stay with the variable-price contract or to switch to a fixed-price contract, depends on the state of the world. Denote $C(D = 1|S = 1, z_0, t)$, $C(D = 1|S = 0, z_0, t)$, $C(D = 0|S = 1, z_0, t)$ and $C(D = 0|S = 0, z_0, t)$ the expected costs of the decision to stay or leave given either state of the world for some individual with a variable-price contract living in a dwelling with characteristics z_0. Assume moreover that decision makers choose the option which minimises their expected cost given the information accumulated so far. Therefore the decision to remain with the variable-price contract is optimal at time t if:

$$\begin{aligned}
&C(D=1|S=1,z_0,t)\Pr[S=1|\mathcal{X}_t,t,z_0,o_0] + C(D=1|S=0,z_0,t)\Pr[S=0|\mathcal{X}_t,t,z_0,o_0] \leq \\
&C(D=0|S=1,z_0,t)\Pr[S=1|\mathcal{X}_t,t,z_0,o_0] + C(D=0|S=0,z_0,t)\Pr[S=0|\mathcal{X}_t,t,z_0,o_0].
\end{aligned}$$

Assuming $C(D = 1|S = 1, z_0, t) < C(D = 0|S = 1, z_0, t)$ and $C(D = 0|S = 0, z_0, t) < C(D = 1|S = 0, z_0, t)$, staying with a variable-price contract is optimal if:

$$\ln \frac{\Pr[S=1|\mathcal{X}_t,t,z_0,o_0]}{\Pr[S=0|\mathcal{X}_t,t,z_0,o_0]} > \ln \frac{C(D=0|S=0,z_0,t) - C(D=1|S=0,z_0,t)}{C(D=1|S=1,z_0,t) - C(D=0|S=1,z_0,t)}.$$

Assume furthermore that, for household i, the log ratio of the cost differences at time t depends on the initial characteristics of the household and of the dwelling, z_{i0}, and on some idiosyncratic unobserved component specific to household i at time t, ϵ_{it}, so that:

$$\ln \frac{C(D=0|S=0,z_0,t) - C(D=1|S=0,z_0,t)}{C(D=1|S=1,z_0,t) - C(D=0|S=1,z_0,t)} = -z_{i0}\frac{\pi_0}{\pi_1} + \epsilon_{it}\frac{1}{\pi_1}. \tag{9}$$

Given all information available at time t, we assume finally that the unobserved component ϵ_{it} is identically and independently distributed across all individuals according to a logistic distribution given χ_t, t, z_{i0}, o_0. Under such conditions, the proportion of individuals in the population who stay with a variable-price contract satisfy the logit specification we set out in Equation (8).

We can provide a set of primitive modeling assumptions on the cost differences which generate precisely this model. Assume that for household i at time t, given the information available then the differences in the conditional expected cost function are of the form:

$$\begin{aligned} C(D=0|S=0,z_0,t)-C(D=1|S=0,z_0,t) &= \exp(z_{i0}\frac{1}{\pi_1}(\phi+\psi_0)+\frac{1}{\pi_1}\eta_{i0t}+u_i) \\ C(D=1|S=1,z_0,t)-C(D=0|S=1,z_0,t) &= \exp(z_{i0}\frac{1}{\pi_1}(\phi+\psi_1)+\frac{1}{\pi_1}\eta_{i1t}+u_i) \end{aligned} \tag{10}$$

where the vector of parameters ϕ and ψ_k, $k=0,1$, are conformable with the dimension of the vector of initial characteristics z_{i0}. u_i is a household specific and time invariant unobserved component. Assuming furthermore that the unobserved components η_{i0t} and η_{i1t} are identically and independently distributed according to an extreme value Gumbel distribution across decisions, states of the world and individuals. The log ratio of the cost differences takes the form $-z_{i0}\frac{1}{\pi_1}(\psi_1-\psi_0)+\frac{1}{\pi_1}(\eta_{i0t}-\eta_{i1t})$ which satisfies Equation (9), and the distributional assumptions imply that the difference $\eta_{i0t}-\eta_{i1t}$ is distributed according to a logistic distribution. We are therefore able to identify $\psi_1-\psi_0$. Finally, we can not recover any information about the vector of parameters ϕ or about the individual specific effect u_i, in other words we are not able to identify the household specific scale factor that is common to both cost differences, i.e. $\exp(z_{i0}\frac{1}{\pi_1}\phi+u_i)$. Therefore in general the analysis of the decision to switch between contract will not be informative about all aspects of the costs functions, but only about the characteristics which determine relative changes in their differences.

Given the initial characteristics z_0, the initial odds o_0 and the specification for λ_t as a function of x_t and t, Equation (8) describes completely the dynamic of the probabilities to switch between contract as new information is revealed. This expression therefore assumes that if the log-odds favours the state of the world where the variable-price contract is preferable to the fixed-price contract, i.e., when $\ln o_t$ is positive, then the probability to remain with the variable-price contract is larger than if the evidence so far supports the state of the world where it is the fixed-price contract that is preferable, i.e., such that $\ln o_t$ is negative.

This specification implies a simple relationship, in terms of log-odds once more, between the proportion of households in the population who keep the variable-price contract and the history of past information:

$$\ln\frac{\Pr[D_t=1|\mathcal{X}_t,t,z_0,o_0]}{\Pr[D_t=0|\mathcal{X}_t,t,z_0,o_0]}=z_0\pi_0+\pi_1\ln\lambda_t+\pi_1\sum_{\tau=1}^{t}\kappa_1^{\tau}\ln\lambda_{t-\tau}+\pi_1\kappa_1^t\ln o_0+\pi_1\kappa_0\frac{1-\kappa_1^t}{1-\kappa_1}, \tag{11}$$

which we can restate as:

$$\ln\frac{\Pr[D_t=1|\mathcal{X}_t,t,z_{i0},o_0]}{\Pr[D_t=0|\mathcal{X}_t,t,z_{i0},o_0]}=\kappa_1\ln\frac{\Pr[D_{t-1}=1|\mathcal{X}_{t-1},t-1,z_{i0},o_0]}{\Pr[D_{t-1}=0|\mathcal{X}_{t-1},t-1,z_{i0},o_0]}+z_{i0}\pi_0(1-\kappa_1)+\pi_1\kappa_0+\pi_1\ln\lambda_t. \tag{12}$$

Hence, the model suggests that the evolution of the log-odds for the proportion of households keeping a variable-price contract is comparable to the dynamic of the posterior log-odds that the fixed-price contract is preferable: the auto-regression parameter is κ_1 in both cases, and the innovation arises in both cases from the information revealed at time t through the likelihood ratio λ_t. Finally, the elasticity of the proportion in the population with a variable-price contract that chooses to change contract at time t, $\eta_{P_0|x_t}$, in response to the log price differential x_t given the accumulated evidence contained in $\mathcal{X}_{t-1}$ takes the form:

$$\begin{aligned}\eta_{\Pr_0|x_t} &\equiv \frac{1}{\Pr[D_t=0|\mathcal{X}_t,t,z_0]}\frac{\partial \Pr[D_t=0|\mathcal{X}_t,t,z_0]}{\partial x_t} \\ &= -\pi_1\eta_{\lambda|x_t}\frac{1}{\Pr[D_t=0|\mathcal{X}_t,t,z_0]}\frac{\partial \Pr[D_t=0|\mathcal{X}_t,t,z_0]}{\partial o_t},\end{aligned} \tag{13}$$

where $\eta_{\lambda|x_t}$ is the elasticity of the likelihood ratio w.r.t. the information x_t. In the absence of unobserved heterogeneity in the initial odds o_0 among households, the expression simplifies to:

$$\eta_{\Pr_0|x_t} = -\pi_1\eta_{\lambda|x_t}\Pr[D_t=1|\mathcal{X}_t,t,z_0].$$

In general, the elasticity depends both on the history of past informations $\mathcal{X}_{t-1}$ and on t for a given value of the current information, x_t. In particular, the model suggest that distinct households will react differently to a realisation of the information at time t if they experienced a distinct history of information and/or if they have kept the variable-price contract for distinct durations.

2.3. Univariate Homoscedastic Information

To operationalise the model, assume that given x_{t-1}, households believe that, conditional on the past and the state of the world, x_t is generated by a autoregressive model with normal innovations so that [6]:

$$\begin{aligned}x_t &\sim \mathcal{N}(\rho_- x_{t-1}+\mu_-,\sigma^2) \text{ if } S=1,\\ x_t &\sim \mathcal{N}(\rho_+ x_{t-1}+\mu_+,\sigma^2) \text{ if } S=0,\end{aligned} \tag{14}$$

such that the mean of the process as well as the parameter of autoregression depends on the state of the world. In our context while it may be sensible to assume that if the state of the world is $S=0$, μ_+ is positive, while if $S=1$ then μ_- is negative this is not required.

Thanks to the normality assumption, the likelihood ratio for the observation of x_t given the history of realisations of the process, $\mathcal{X}_{t-1}$, takes the form:

$$\begin{aligned}\ln\lambda_t \equiv \ln\frac{P[x_t|S=1,\mathcal{X}_{t-1},t]}{P[x_t|S=0,\mathcal{X}_{t-1},t]} &= -\frac{1}{2\sigma^2}\left((x_t-\rho_- x_{t-1}-\mu_-)^2-(x_t-\rho_+ x_{t-1}-\mu_+)^2\right)\\ &= -\frac{1}{\sigma^2}(\Delta\rho\; x_{t-1}+\Delta\mu)(x_t-\bar{\rho}x_{t-1}-\bar{\mu}),\end{aligned} \tag{15}$$

[6] Hence the model we develop here is not assuming that households necessarily consider or believe in the "true" model that generates the data. A more advanced theory would integrate a search through possible models.

where $\Delta\mu \equiv \mu_+ - \mu_-$, $\Delta\rho \equiv \rho_+ - \rho_-$ and $\bar{\mu} \equiv \frac{1}{2}(\mu_+ + \mu_-)$.

The recent evidence, x_t, favours $S_t = 1$ (in the sense that it increases the posterior odds that $S_t = 1$) if i) x_t is small enough when $\Delta\rho\ x_{t-1} + \Delta\mu > 0$, i.e., $x_t \leq \rho x_{t-1} + \bar{\mu}$, since in that case $\ln \lambda_t \geq 0$; or ii) if x_t is larger than $\rho x_{t-1} + \bar{\mu}$, whenever $\Delta\rho\ x_{t-1} + \Delta\mu < 0$. Although the log likelihood ratio response to the current realisation of the information depends on the past realisations, the weight of evidence in favour of staying with the variable-price contract increases as long as $\ln \lambda_t$ is positive. Figure 2 illustrates how the sign of the likelihood ratio vary across the (x_{t-1}, x_t) plane for two distinct configurations of the parameters. Observe that the regions of the plane (x_{t-1}, x_t) such that the log likelihood ratio is positive, i.e., the current observation supports the state of the world where the variable price is cheaper, can vary markedly depending on the parameter values.

Figure 2 Sign of Log Likelihood Ratio, eq. (15)

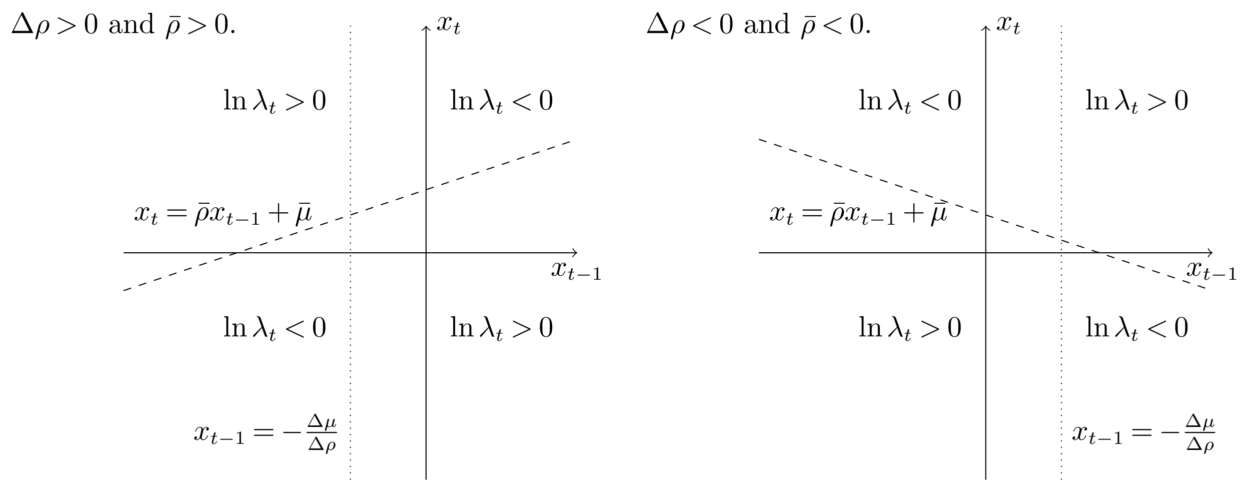

Note: The two figures show how the sign of the log likelihood ratio, given in (15), varies across regions of the (x_{t-1}, x_t) plane. The dotted and the dashed lines are the limits of the regions. In both figures, on the left [resp. right] of the dotted line the elasticity $\eta_{\lambda|x_t}$ is negative for all values of x_t if $\Delta\rho > 0$ [resp. $\Delta\rho < 0$].

The expression of the elasticity of the likelihood ratio w.r.t x_t [7] provides an alternative understanding of the effect of the current information on the probability to stay on a variable-price contract:

$$\eta_{\lambda|x_t} = -\frac{1}{\sigma^2}(\Delta\rho\ x_{t-1} + \Delta\mu), \tag{16}$$

Hence, the elasticity of the likelihood ratio w.r.t. x_t is positive if $\Delta\rho\ x_{t-1} + \Delta\mu$ is negative, and it is positive otherwise. The response of the likelihood ratio is independent of the current information: it increases whenever $\Delta\rho x_{t-1} < -\Delta\mu$ and decreases otherwise. Moreover, it is possible to find a

[7] We think of x_t as the log price differential.

couple of observations (x_{t-1}, x_t) for which the elasticity $\eta_{\lambda|x_t}$ is negative and the log-likelihood ratio is positive.

Finally, in the special case where $\Delta\rho = 0$ and $\Delta\mu > 0$, where the autoregressive parameter is identical in the two states of the world and the price differential is negative when the flexible price is preferable and positive otherwise, the log likelihood ratio decreases with larger values of x_t. That is, the larger the value of x_t, the larger the evidence against the state where the variable price is preferable to the fixed price on average. Clearly, our analysis shows that this behaviour is the general case even in the homoscedastic case.

2.4. Univariate Heteroscedastic Information

We can extend the previous specification and assume that households believe that the conditional variances of the innovation depend on the state of the world:

$$\begin{aligned} x_t &\sim \mathcal{N}(\rho_- x_{t-1} + \mu_-, \sigma_-^2) \text{ if } S = 1, \\ x_t &\sim \mathcal{N}(\rho_+ x_{t-1} + \mu_+, \sigma_+^2) \text{ if } S = 0. \end{aligned} \tag{17}$$

Figure 3 Sign of Log Likelihood Ratio, eq. (18)

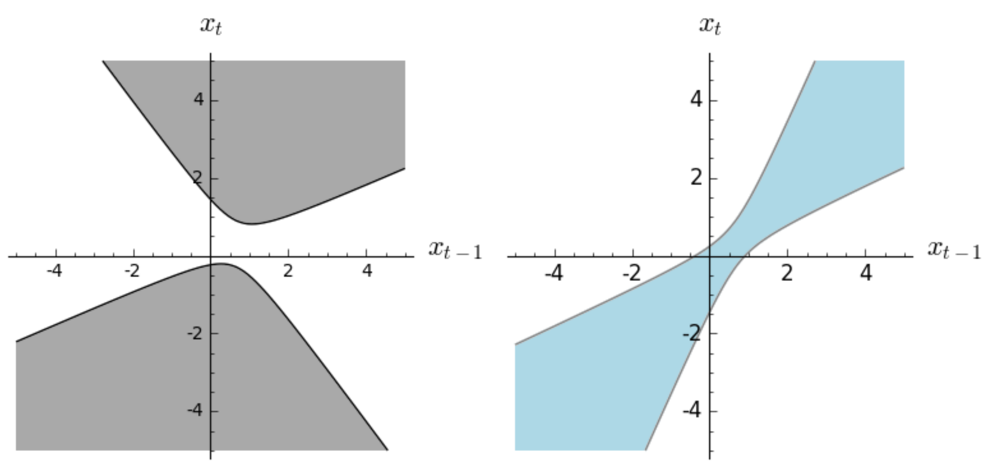

Note: The two figures show how the sign of the log likelihood ratio, given in (18), can vary across regions of the (x_{t-1}, x_t) plane for distinct values of the parameters. The shaded area in each case indicates the region in the plane where the log likelihood ratio is positive, i.e. (x_{t-1}, x_t) increases the posterior odds in favour of $S = 1$. The parameters for the regions on the left hand pane are: $\mu_- = -0.1$, $\mu_+ = 0.1, \rho_- = 0.6, \rho_+ = 0.3, \sigma_- = 0.65$ and $\sigma_+ = 0.55$. To produce the right hand pane, the parameters are identical except for $\sigma_- = 0.55$ and $\sigma_+ = 0.65$.

The log likelihood ratio is:

$$\ln \lambda_t = \ln(\frac{\sigma_+}{\sigma_-}) - \frac{1}{2}\Big(\frac{1}{\sigma_-^2}(x_t - \rho_- x_{t-1} - \mu_-)^2 - \frac{1}{\sigma_+^2}(x_t - \rho_+ x_{t-1} - \mu_+)^2\Big),$$

which can be expanded into the following expression for the log likelihood ratio at time t:

$$\ln\lambda_t = \ln(\frac{\sigma_+}{\sigma_-}) + \frac{1}{2}\Delta[\frac{\mu^2}{\sigma^2}] + \frac{1}{2}\Delta[\frac{1}{\sigma^2}]x_t^2 + \frac{1}{2}\Delta[\frac{\rho^2}{\sigma^2}]x_{t-1}^2 - \Delta[\frac{\rho}{\sigma^2}]x_t x_{t-1} - \Delta[\frac{\mu}{\sigma^2}]x_t + \Delta[\frac{\rho\mu}{\sigma^2}]x_{t-1} \quad (18)$$

where $\Delta[\frac{\mu^j\rho^k}{\sigma^2}] \equiv \mu_+^j\rho_+^k\sigma_+^{-2} - \mu_-^j\rho_-^k\sigma_-^{-2}$ for $j,k = 0,1,2$. Observe moreover that while $\Delta[\frac{\mu}{\sigma^2}]$ is always positive if we assume $\mu_+ > 0 > \mu_-$, the signs of the other terms $\Delta[\frac{\mu^j\rho^k}{\sigma^2}] \equiv \mu_+^j\rho_+^k\sigma_+^{-2} - \mu_-^j\rho_-^k\sigma_-^{-2}$ for $j,k = 0,1,2$ are unrestricted in general.

The range of values of x such that likelihood ratio is positive, i.e., the range of values which increase the posterior odds of $S = 1$, depends now on the sign of $\Delta[\sigma^{-2}]$. If $\Delta[\sigma^{-2}]$ is negative, $\sigma_- < \sigma_+$, the range of values of x_t consistent with a positive log-likelihood ratio takes the form of a non empty[8] interval $[\underline{x}(x_{t-1}), \bar{x}(x_{t-1})]$, where $\underline{x}(x_{t-1})$ and $\bar{x}(x_{t-1})$ are the values of x which set λ_t to zero. If instead $\Delta[\sigma^{-2}]$ is positive, $\sigma_+ < \sigma_-$, the values of x consistent with a positive log likelihood ratio, i.e., the evidence supports $S = 1$, is the union of two intervals $[-\infty, \underline{x}(x_{t-1})] \bigcup [\bar{x}(x_{t-1}), +\infty]$.[9]

The log-likelihood ratio at time t, and therefore the sum of log-likelihood ratios, does no longer respond to x_t in a monotone fashion since:

$$\frac{\partial \ln\lambda_t}{\partial x_t} = \frac{\partial \sum_{\tau=1}^{t} \ln\lambda_\tau}{\partial x_t} = \Delta[\frac{1}{\sigma^2}]x_t - \Delta[\frac{\rho}{\sigma^2}]x_{t-1} - \Delta[\frac{\mu}{\sigma^2}]. \quad (19)$$

If $\Delta[\frac{1}{\sigma^2}] > 0$, that is, households believe that the price differential is more variable when the average price differential is positive, the log likelihood ratio increases whenever $x_t > \Delta[\frac{\rho}{\sigma^2}]/\Delta[\sigma^{-2}]x_{t-1} + \Delta[\frac{\mu}{\sigma^2}]/\Delta[\sigma^{-2}]$. Hence, in this case, only a large enough x_t increases the log-likelihood ratio and therefore reduces the probability to leave the variable-price contract. If instead $\Delta[\sigma^{-2}] < 0$, values of x_t smaller than $\Delta[\frac{\rho}{\sigma^2}]/\Delta[\sigma^{-2}]x_{t-1} + \Delta[\frac{\mu}{\sigma^2}]/\Delta[\sigma^{-2}]$ lead to larger log-likelihood ratio. Clearly, if $\sigma_- = \sigma_+$ so that $\Delta[\sigma^{-2}] = 0$, we recover the results obtained in the homoscedastic case.

Figure 4 describes two examples of the behaviour of the log odds $\ln o_2$, assuming that the log likelihood ratio takes the form given in equation (2.4), after three realisations of x_t for two distinct parametrisations. The first two realisations are set at $x_0 = 0$ and $x_1 = -1$ while x_2 is allowed to vary over the range $(-3, 2)$. The figure illustrates the fact that the log odds can be positive (or negative) and yet display either a positive or a negative slope relative to x_2. For $x_2 = -3$, $\ln o_2$ is positive for both parametrisations, but in one case the log odds is an increasing function of x_2 at that point whereas in the other case the log odds is a decreasing function of x_2. Similarly, for $x_2 = 1.5$, both log odds are negative, and in one case the log odds is increasing at that point while in the other case it is decreasing.

[8]When $\sigma_- < \sigma_+$, the quantity $\ln(\frac{\sigma_+}{\sigma_-}) + \frac{1}{2}\Delta[\frac{\mu^2}{\sigma^2}]$ is always positive.

[9]The intersection between the two intervals is empty if $\ln(\frac{\sigma_+}{\sigma_-}) + \frac{1}{2}\Delta[\frac{\mu^2}{\sigma^2}]$ is negative.

Hence, a new realisation of x_t adds to evidence, in the form of $\ln \lambda_t$, to the prior log odds $\kappa_0 + \kappa_1 \ln o_{t-1}$. For some values of x_t, the added information increases the log odds. It is likely, however, that there are alternative values such that the log-likelihood ratio, and therefore the log odds, can reach even higher levels. If such a value is to the right of the realised x_t , then the log odds will be increasing with x_t. If, on the other hand, it is to the left of the realised value, the log odds will be decreasing. However, the probability to keep a variable-price contract increases whenever a new realisation of x_t increases the log odds, $\ln o_t$, which arises if and only if $\ln \lambda_t > 0$. As we have seen in our analysis above, realisations of x_t such that $\ln \lambda_t > 0$ can take place whether $\eta_{\lambda|x} > 0$ or $\eta_{\lambda|x} < 0$.

Figure 4 Log Odds, $\ln o_2$ for $x_0 = 0$, $x_1 = -1$, eq. (4)

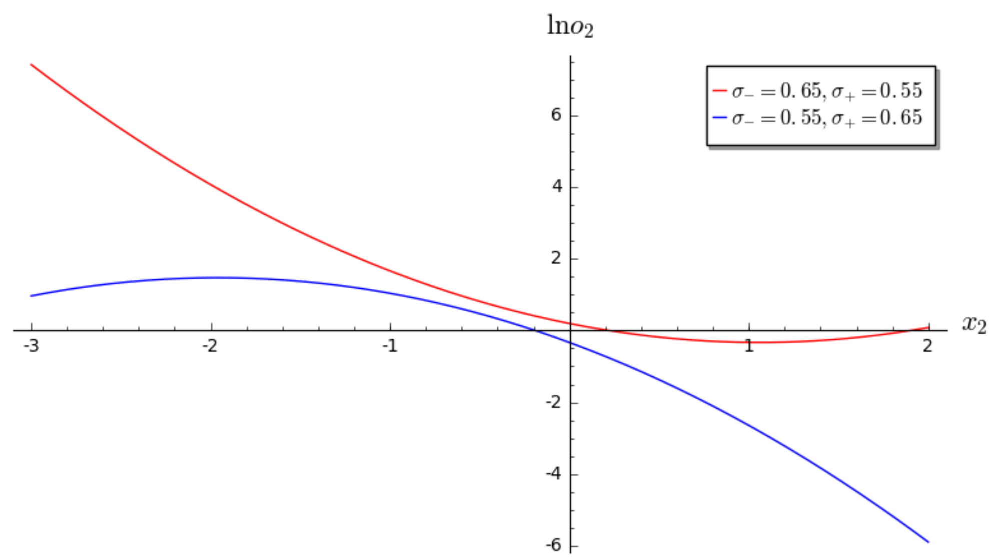

Note: Assuming that the first two realisations are $x_0 = 0$ and $x_1 = -1$, the figure shows how the log odds $\ln o_2$, changes with x_2. The other parameters to define the two curves are: $\mu_- = -0.1,\ \mu_+ = 0.1,\ \rho_- = 0.6$, and $\rho_+ = 0.3$.

2.5. Bivariate information

Assume now that the price differential and the local temperature deviations from long term trends are observed by the decision maker. x_t is now multivariate (here bivariate). Assume that the first component, x_{1t} of x_t, measures the local temperature deviations from long run averages while the second component, x_{2t}, measures the log price differential between the variable-price and the fixed-price contracts.

Because our data covers a short period of time, the households are assumed to believe that the process which describes the evolution of the local climate is independent of the state of the world which describes the evolution of the log price differential.

We assume further that given the state of the world (S) and given the history $\mathcal{X}_{t-1}$ until $t-1$, we can decompose the likelihood $P[x_t|S,\mathcal{X}_{t-1},t]$ in terms of the product of the marginal density for the temperature information contained in x_{1t}, $f_1(x_{1t}|x_{1,t-1})$, and of the conditional density of the log price differential, $f_2(x_{2t}|S,x_{2,t-1},x_{1,t},x_{1,t-1})$:

$$\Pr[x_t|S,\mathcal{X}_{t-1},t] = f_1(x_{1t}|x_{1,t-1})f_2(x_{2t}|S,x_{2,t-1},x_{1,t},x_{1,t-1}). \tag{20}$$

The log-likelihood ratio for the observation of x_t at time t then takes the simpler form:

$$\begin{aligned}\ln\lambda_t &= \frac{\Pr[x_t|S=1,\mathcal{X}_{t-1},t]}{\Pr[x_t|S=0,\mathcal{X}_{t-1},t]}\\ &= \frac{f_1(x_{1t}|x_{1,t-1})f_2(x_{2t}|S=1,x_{2,t-1},x_{1,t},x_{1,t-1})}{f_1(x_{1t}|x_{1,t-1})f_2(x_{2t}|S=0,x_{2,t-1},x_{1,t},x_{1,t-1})} = \frac{f_2(x_{2t}|S=1,x_{2,t-1},x_{1,t},x_{1,t-1})}{f_2(x_{2t}|S=0,x_{2,t-1},x_{1,t},x_{1,t-1})},\end{aligned} \tag{21}$$

which depends only on the conditional densities given each state of the world of the log price differential. Obviously, this does not mean that the history of the local temperature at time t is irrelevant. The model suggests that it matters in so far as it is able to help predicting current prices and therefore enters the conditional densities, $f_2(x_{2t}|S=1,x_{2,t-1},x_{1,t},x_{1,t-1})$.

To extend the model specification we discussed in the univariate case, assume that the bivariate process $X_t \equiv (x_{1t},x_{2t})'$ is described by a state contingent triangular VAR(1) model of the form:

$$\begin{aligned} x_{1t} &= \nu_0 + \psi x_{1,t-1} + \epsilon_{1t},\\ x_{2t} &= \rho_{1S}x_{1,t-1} + \rho_{2S}x_{2,t-1} + \mu_S + \epsilon_{2t},\end{aligned} \tag{22}$$

with $\mu_1 \equiv \mu_-$, $\mu_0 \equiv \mu_+$, $\rho_{k1} \equiv \rho_{k-}$, $\rho_{k0} \equiv \rho_{k+}$, for $k=1,2$. We assume that the innovations $\epsilon_t \equiv (\epsilon_{1t},\epsilon_{2t})'$ are conditionally normally distributed given the past information, $\mathcal{X}_{t-1}$. Hence, given the information about the realisation of the temperature x_{1t} at time t, the distribution of x_{2t} remains normal. However, in general the parameters which define the conditional distribution will depend on the state. Consider the form of the variance covariance of the innovation given the state of the world and the past information:

$$\mathbb{V}[\epsilon_t|S,\mathcal{X}_{t-1}] = \begin{bmatrix} \sigma_1^2 & \omega_S\sigma_1\sigma_{2,S}\\ \omega_S\sigma_1\sigma_{2,S} & \sigma_{2,S}^2\end{bmatrix}, \tag{23}$$

where $\sigma_{2,S}^2$ is the variance of the innovation in the definition of x_{2t} given the state of the world. ω_S is the correlation between the innovation given the state of the world. We set $\sigma_{21} \equiv \sigma_{2-}$, $\sigma_{20} \equiv \sigma_{2+}$, $\omega_1 \equiv \omega_-$, $\omega_0 \equiv \omega_+$. In general, we assume $\sigma_{2-} \neq \sigma_{2+}$ and $\omega_- \neq \omega_+$. This specification guarantees that the variance of x_{1t} is independent of the belief of the state of the world, while the conditional expectation and the conditional variance of ϵ_{2t} given ϵ_{1t} will depend on the belief of the state of the world:

$$\mathbb{E}[\epsilon_{2t}|S,\mathcal{X}_{t-1},\epsilon_{1t}] = \omega_S\sigma_{2S}\epsilon_{1t}, \tag{24}$$

$$\mathbb{V}[\epsilon_{2t}|S,\mathcal{X}_{t-1},\epsilon_{1t}] = \sigma^2_{2S}(1-\omega^2_S) \equiv \tilde{\sigma}^2_{2S}. \tag{25}$$

Given the state of the world, the history $\mathcal{X}_{t-1}$ and x_{1t}, the conditional expectation of x_{2t} takes the form:

$$\begin{aligned} x_{2t} &= \rho_{1S}x_{1,t-1} + \rho_{2S}x_{2,t-1} + \mu_S + \frac{\omega_S\sigma_{2S}}{\sigma_1}(x_{1t} - \nu_0 - \psi x_{1,t-1}) + \tilde{\epsilon}_{2t} \\ &= \frac{\omega_S\sigma_{2S}}{\sigma_1}x_{1t} + (\rho_{1S} - \frac{\omega_S\sigma_{2S}}{\sigma_1}\psi)x_{1,t-1} + \rho_{2S}x_{2,t-1} + \mu_S - \frac{\omega_S\sigma_{2S}}{\sigma_1}\nu_0 + \tilde{\epsilon}_{2t} \\ &= \tilde{\rho}_{0S}x_{1t} + \tilde{\rho}_{1S}x_{1,t-1} + \rho_{2S}x_{2,t-1} + \tilde{\mu}_S + \tilde{\epsilon}_{2t} \end{aligned} \tag{26}$$

where $\tilde{\epsilon}_{2t}$ is a normally distributed innovation with conditional variance equal to $\mathbb{V}[\epsilon_{2t}|S,\mathcal{X}_{t-1},\epsilon_{1t}]$. The last expression collects the composite terms into state dependent parameters (denoted by ˜). The above equation shows that even in the normal case the conditional densities $f_2(x_{2t}|S, x_{2,t-1}, x_{1,t}, x_{1,t-1})$, for $S = 0, 1$, are described in general by parameters which depend on the state of the world (if $\sigma_{20}\omega_0 = \sigma_{21}\omega_1$, the conditional mean is state invariant, although the conditional variance may still be state dependent).

The log likelihood ratio takes the form:

$$\begin{aligned} \ln\lambda_t = \ln(\frac{\tilde{\sigma}_{2+}}{\tilde{\sigma}_{2-}}) - \frac{1}{2}\Big(&\frac{1}{\sigma^2_{2-}}(x_{2t} - \tilde{\rho}_{0-}x_{1t} - \tilde{\rho}_{1-}x_{1,t-1} - \rho_{2-}x_{2,t-1} - \tilde{\mu}_-)^2 - \\ &\frac{1}{\sigma^2_{2+}}(x_{2t} - \tilde{\rho}_{0+}x_{1t} - \tilde{\rho}_{1+}x_{1,t-1} - \rho_{2+}x_{2,t-1} - \tilde{\mu}_+)^2\Big) \end{aligned} \tag{27}$$

Adding additional information, here the information about temperature x_{1t}, modifies the conditional mean for the price differential x_{2t} without modifying significantly the conclusions reached in Section 2.3 or 2.4. Further algebra (of the kind we report earlier) will give expressions for the log-likelihood ratio similar to the one we discuss in Equation (18). In particular, we can show that there are regions of the space $(x_{1t-1}, x_{2t-1}, x_{1t}, x_{2t})$ such that the log likelihood ratio is positive, i.e., that supports the belief that the state of the world is $S = 1$, and that over such region the log-likelihood ratio can be increasing or decreasing with the current price differential observation x_{2t}.

3. Data

The data used in this paper originates from the customer database of Skellefteå Kraft, one of the bigger electricity retailers in Sweden (see `www.skekraft.se`). The source material has been anonymized and identification of households is not possible. We use a sub-set of this data, consisting of 3353 new customers that have started a variable-price contract during the period June 2010 to February 2012 and did not have an electricity contract with Skellefte Kraft prior to this. Although we do not observe what contract (if any) these households had prior to the sample period, we do know that the choice of variable-price contract with Skellefte Kraft was an active choice (e.g., compared to default contracts which households are automatically assigned to if they do not make an active choice; see Section 1). The households in the sample have either remained on their variable-price contract throughout the sample period which ends in June 2014, or transitioned to either a fixed-price contract with the current retailer, or transitioned to another retailer (in which case we do not observe to what contract type they have switched).

Most of the households in the sample (76 percent) are located in the same postal area as the retailer in northern Sweden (Skellefteå), whereas the Swedish population is concentrated in the middle and south of Sweden. The location of the households in the sample is expected, because one important determinant of retailer choice is assumed to be geographical location and closeness to the retailer. For example, Goett et al. (2000), Revelt and Train (2000) and Yang (2014) demonstrate that households prefer their local company to any other retailer. However, the sample contains observations on households from all 21 postal areas in Sweden. For example, roughly 8 percent of the households live in Stockholm, and another 3 percent in Gothenburg.

With data on households from only one supplier out of many in Sweden, there is the obvious risk that the data used in this paper is not representative of the Swedish population, and that the results cannot readily be generalized to households choosing between electricity contracts from other suppliers. However, the Swedish electricity retail market is typically thought of as a competitive market, and retailers are expected to offer similar contract terms (e.g., Damsgaard et al. (2005) and Brännlund et al. (2012)). In that sense, the decision to remain on a variable-price contract or switch to a fixed-price contract within Skellefteå Kraft should be relatively similar to the corresponding decision within other retailers.

The data includes detailed household-level information about the current electricity contracts, prices per kWh, electricity consumption in kWh, geographical location and housing type (villa or flat). Unfortunately, the customer database lacks any household-level information about income and other household characteristics and only includes socio-economic information at the zip-code level using census data with figures from 2014 for annual per-capita income, family size, age and education level.[10] For the whole of Sweden, the average number of persons in each zip-code is 918, allowing for precise neighbourhood matching with census data. Matching individual-level data with zip code-level data is frequently used in economics, including residential electricity demand (Borenstein (2012); Ito (2014)). Temperature for each postal area originates from SMHI (Swedish Meteorological and Hydrological Institute; see `www.smhi.se`). Both price and temperature data is

[10] Specifically, income and age are coded as dummy variables for income and age groups (10 groups each based in income and age distributions), family size is a dummy variable taking the value one if more than 30 percent of the households have two or more inhabitants, and education is defined as a one if more than 40 percent of the people in the zip-code has a university degree

Table 1 Summary Statistics

	Mean	Standard deviation
Variable price, öre/kWh	43.35	11.21
Fixed price, öre/kWh	49.46	0.72
Temperature (Celsius)	4.58	8.18
Monthly kWh, electrically heated villas	1473.13	1020.00
Monthly kWh, non-electrically heated villa	964.92	869.40
Monthly kWh, flats	735.01	843.33
Income (1000's SEK)	273.52	55.85
Age	49.07	4.55
Flats, share	0.79	
Non-electrically heated villas, share	0.12	
Electrically heated villas, share	0.09	
Family size, share (more than two persons = 1	0.34	
University, share (University degree = 1)	0.22	
Number of Households in sample	3353	
Average Number of Month in sample	18	

Note: Prices are excluding transmission, green certificate and taxes, which are independent of contract choice.

seasonally adjusted (by postal area in the case of temperature), and since price and temperature data is available from 2006, the seasonally adjusted variables may be thought of as deviations from the long-run monthly average price and temperature. Descriptive statistics of the variables used in the empirical application are found in Table 1, and the seasonally adjusted monthly temperature for Stockholm and Umeå is illustrated in Figure 5.

The variable price per kWh is lower than the corresponding fixed price, as expected, because households on the latter type of contract pays a premium to avoid monthly price variability: the average price difference is six öre/kWh. In addition to the price per kWh, households pay for transmission, green certificate and taxes, which are independent of contract choice. These costs amounts to roughly 50 öre/kWh, plus a fixed annual transmission fee varying between 1500 and 15000 SEK depending on the required amperage.

Roughly 80 percent of the households in the sample are living in flats. For the Swedish population, the corresponding figure is less than 50 percent. One explanation to the large share of flats in the sample could be that variable-price contracts may be more common for flats. Electricity consumption is highest for electrically heated villas, as expected, and smaller for non-electrically heated villas and flats, and the electricity usage levels are comparable to those reported by, e.g., the Swedish Energy Agency (see `www.energimyndigheten.se/en`). The relatively large standard deviations illustrate variation in electricity usage across seasons, with average electricity demand during winter being substantially higher than during summer. Since temperature, income, age, family size and university are zip-code level averages, these variables obviously corresponds to population figures.

Figure 5 Seasonally adjusted temperature for Stockholm and Umeå

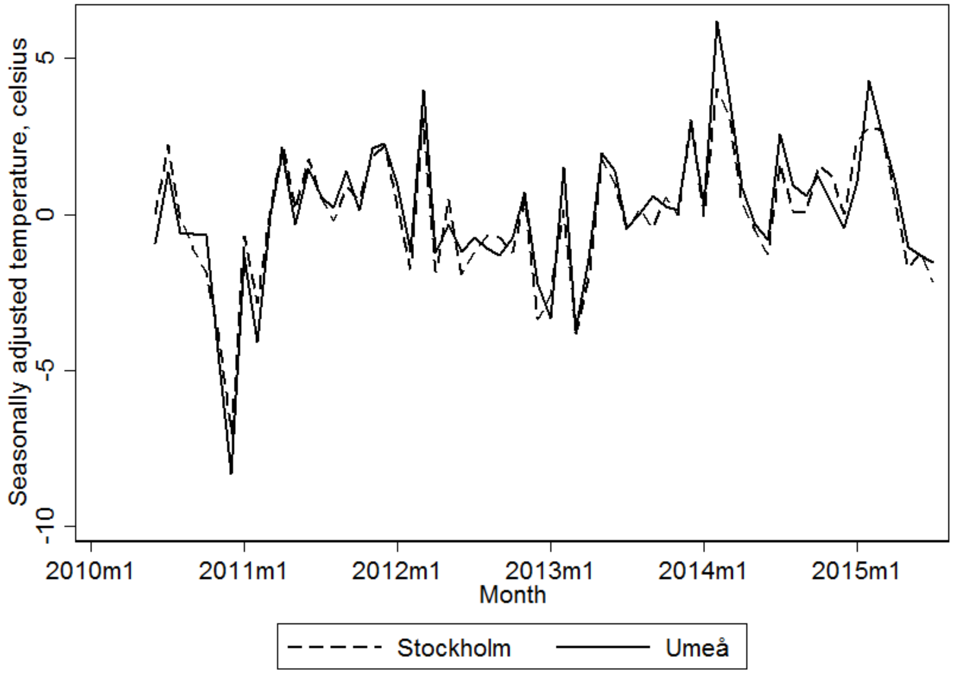

Note: The figure shows seasonally adjusted temperature (in degrees celsius, from month average) for Stockholm, located in the centre of Sweden and Umeå, located in the north of Sweden.

Relating to the previous discussion on electricity demand being similar before and after a switch, we compare the percentage change before and after a switch to the percentage change between any other months in our data, and cannot reject the null hypothesis that the percentage changes are equal. This confirms the idea that consumption should be more or less identical before and after a switch. However, we do find a significant difference in percentage change in costs (price times quantity) before and after a switch compared to the percentage change between any other months.

730 of the households in the sample have transitioned to a fixed-price contract with the current retailer[11] (in the Swedish population, between 10 and 20 percent of the households switch between contracts and/or retailer each year (SCB (2014)). As is evident from Figure 6, most of these switches occurred during 2010 and 2011, with virtually no transitions during 2012. The number of switches peaked during October 2010 with roughly 100 switches, corresponding to roughly seven percent of the households in the data that month. Note that time periods with many transitions occur some months after positive price differentials, and that periods of negative price differentials are associated with few or zero transitions, as would be expected. Also note that the number of transitions to another retailer ("Transitions to unknown") is less variable over time, but that there

[11] $D = 0$ only when a transition from a variable-price contract to a fixed-price contract (with Skellefte Kraft) is observed.

Figure 6 Seasonally adjusted log price differential and number of transitions per month

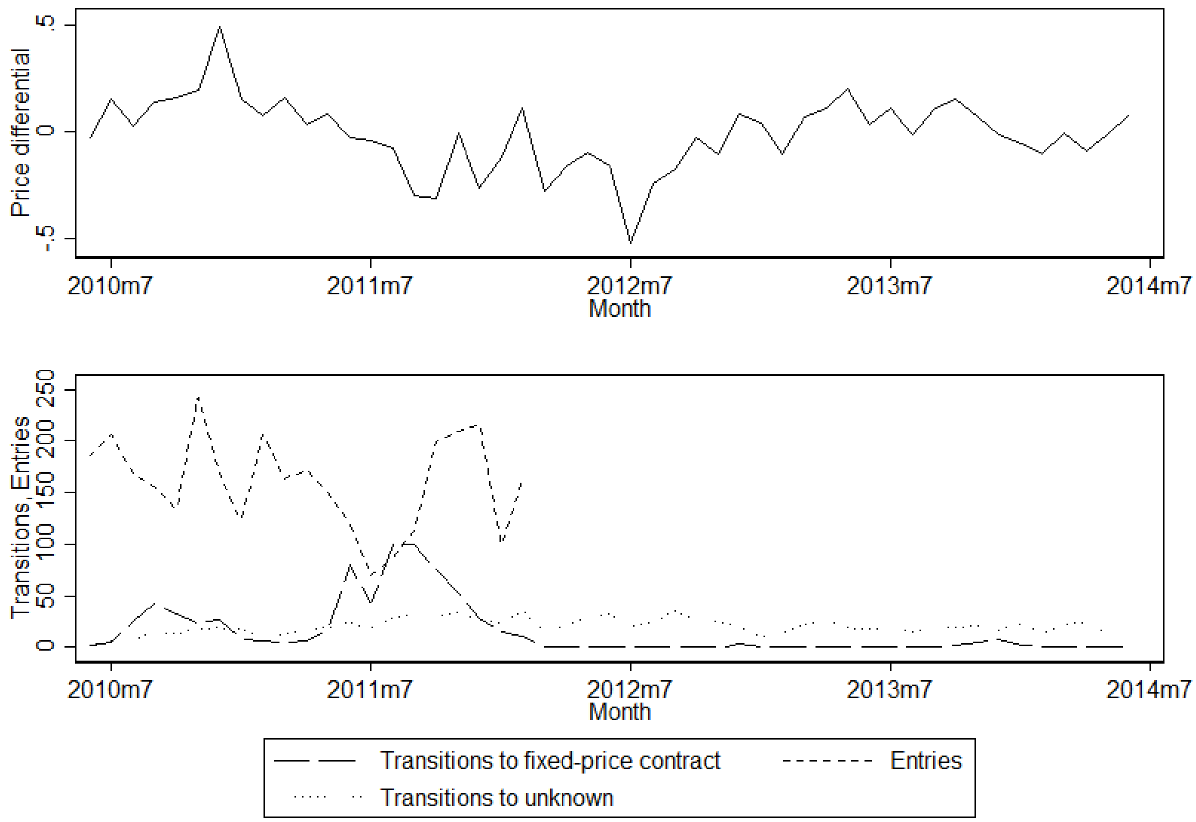

Note: "Transitions to fixed-price contract" refers to transitions within the same supplier to a fixed-price contract, "Transitions to unknown" refers to transitions to another retailer, and "Entries" refers to number of new variable-price contracts per month. Our sample only includes households starting a variable-price contract between June 2010 and February 2012, so number of entries in the sample are zero after February 2012. "Price differential" is the seasonally adjusted log price differential between the variable price and the fixed price.

are somewhat more such transitions after periods of positive price differentials compared to periods of negative price differentials.

The proportion of households remaining on variable-price contracts over time is illustrated in Figure 7, where we illustrate the survival function for households on variable-price contracts using the Kaplan-Meier estimator. Evidently, the survival function drops rather quickly, but becomes rather flat after roughly 15 months. Comparing with the number of transitions per month in Figure 6, we observe many transitions in the beginning of the sample, where the households in the sample have spent a relatively short time on the variable-price contract. Obviously, the difference between the survival function depicted in Figure 7 and the model outlined in this paper is that the latter allow us to condition the survival probability on distinct histories of information, thereby allowing

Figure 7 Kaplan-Meier survival estimate

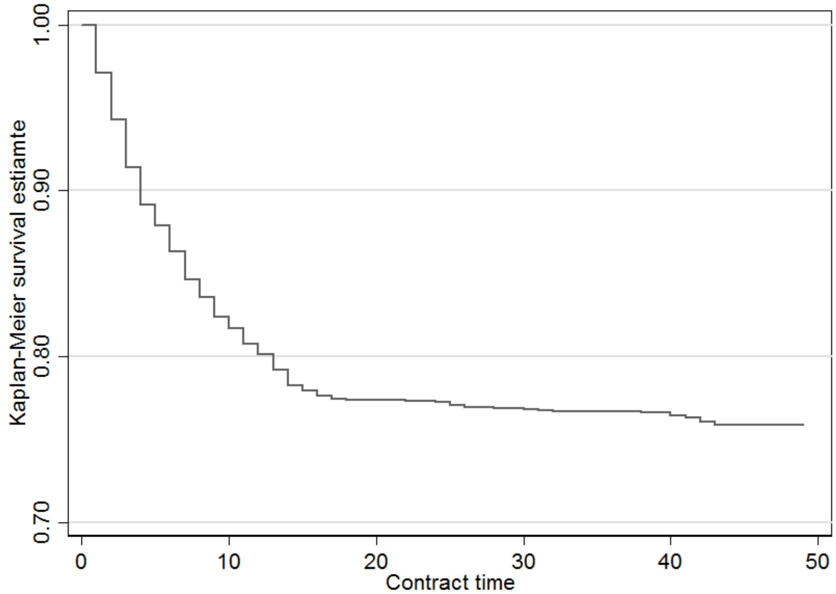

Note: The figure shows the Kaplan-Meier estimator for the proportion of households remaining on variable-price contracts over time.

us to understand the sensitivity of the probability of remaining on a variable-price contract to variation in the price and temperature histories.

Finally, the number of new variable-price contracts per month ("Entries") varies between 70 and 243. Since households start their variable-price contract at different points in time, they face different histories of prices. This variation in price histories, together with the variation in temperature across location, insures the identification of the parameters of the model.

4. Estimation and Results

4.1. Estimation

For each household i we observe a history of choices, $\mathcal{D}_t$, a history of the realisation of the variables measuring the information available to decision makers up to time T_i, $\mathcal{X}_t$, as well as the variable describing the household initial conditions, z_0, i.e., we observe $(\mathcal{D}_{T_i}, \mathcal{X}_{T_i}, z_0)$. The specification of the probabilities we provided earlier allow for a simple expressions to the individual contributions to the (conditional) likelihood:

$$\mathcal{L}_i((\mathcal{D}_{T_i}, \mathcal{X}_{T_i}), \theta) = \prod_{t=1}^{T_i} \Pr[D_{it} = 1 | \mathcal{X}_t, t, z_{i0}]^{D_{it}} \Pr[D_{it} = 0 | \mathcal{X}_t, t, z_{i0}]^{1-D_{it}}, \tag{28}$$

where the vector θ collects all the parameters of the model. From this expression, calculating the sample log-likelihood is straightforward. The information vector describing the initial conditions, z_{i0}, includes building type (electrically heated villa, non-electrically heated villa or flat), income level, family size (more than two persons or not), education (university degree or not) and age. The log price price differential (log of variable price divided by fixed price) and temperature information is defined relative to their monthly averages (at the postal area level for temperature).

The normality assumption of the conditional distribution in either state of the world which we focus on is useful since it leads to expressions for the log likelihood ratio which are linear in the statistics of the data. Hence the log odds $\ln o_t$, which is a (weighted) sum of log likelihood ratios, is itself a linear form of the statistics of the data x_{1t}, x_{2t} and their squares $x_{1t}^2, x_{1t-1}^2, x_{2t}^2, x_{2t-1}^2$, and cross products $x_{1t}x_{2t}, x_{1t}x_{1t-1}, x_{2t}x_{2t-1}, x_{1t}x_{2t-1}, x_{1t-1}x_{2t}$ etc... The distributional parameters which characterise the conditional distribution in each state of the world determine the function of the parameters, i.e. of the form $\Delta[\frac{\mu^j \rho_l^k}{\sigma^2}]$, for each particular statistics that enters the log odds. The scale of these parameters itself is in general not identified since the contribution of the log odds to the probability to keep a flexible price contract is scaled by π_1. In our applied work we do not attempt to recover the distributional parameters from the estimated parameters. Hence the parameter θ includes as well the parameters of the initial characteristics and the parameter of the trend and of the information discounting.

Table 2 summarizes estimation results for the different specifications, including the number of parameters, the estimated trend and discount, the log-likelihood, AIC and BIC. The complete list of parameter estimates for all specifications can be found in the Appendix in Tables 3 and 4. Evidently, both the log-likelihood and the two information criterion suggest a better fit for the heteroskedastic specifications than for the homoskedastic specification, and that including trend and discount further improves the model fit in all cases. For each specification Table 2 reports the estimated values for the trend and the discount parameters. These parameters are precisely estimated in all cases. While the discount rate is measured close to unity in all cases, the parameter estimates for the trend vary between specifications and is negative with the most complex specification. In a first approximation we interpret the parameter of the trend as the parameter κ_0 in Equation (5).[12] When κ_0 is negative, the decision maker scales down the posterior odds she reached at the end of period t towards zero to produce her prior odds at the start of period $t+1$. If we interpret the parameter estimate in the last column as $\kappa_0 = -0.467$ this amounts to scaling down the posterior odds by 0.63. Together with the estimated discount rate, the results for the last specification suggest that decision makers forget past evidence and systematically downplay the information in favour of variable prices by a considerable factor.

[12]Note that in general the log likelihood ratio contains a constant term which will accumulate in the log odds and yield a time dependent function which is identical the last term of Equation (5). It is not possible to identify separately this constant term from κ_0.

Table 2 Estimation Summary

Specification	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
No. of param.	2	2	2	3	3	3	14	14	14
Init. cond.	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Trend		0.126***	0.152***		0.0708***	0.107***		-0.331***	-0.467***
		(0.007)	(0.013)		(0.008)	(0.013)		(0.040)	(0.051)
Discount			0.980***			0.978***			0.894***
			(0.007)			(0.005)			(0.015)
Number obs.	83786	83786	83786	83786	83786	83786	83786	83786	83786
Log-Likelihood	-3857.7	-3624.4	-3620.3	-3595.5	-3548.9	-3538.9	-3398.8	-3360.1	-3323.9
AIC	7765.3	7300.7	7294.7	7243.0	7151.8	7133.9	6867.6	6796.2	6725.8
BIC	7998.7	7543.4	7546.8	7485.7	7403.9	7395.3	7213.0	7151.0	7089.9

Note: i) The table report maximum likelihood estimates based on the maximisation of expression (28) on the information in the sample. All specification accounts for the same initial characteristics. The number of parameters in the model of the log-likelihood ratio is reported in the second row.

ii) Columns 1-3 report on the estimation results for the univariate homoskedastic model of the log likelihood ratio (described in Equation 15) where the information is the log price differential only. Column 4-6 report results for the univariate homoskedastic model (described in Equation 15) where the mean log price differential at time t depends on the local temperature deviation from monthly means. Column 7-9 report results for the bivariate heteroskedastic model (described in Equation 27) where the information is both the log price differential and the temperature deviations from monthly averages. We report separately the estimates of the trend and discount parameters and their standard errors (in parentheses).

iii) The number of observations refer to the number of individuals × months observed in the sample.

iv) The log-likelihood statistics, Akaike's and Schwarz's Bayesian information criteria statistics are reported for each model. While a larger log-likelihood suggests a better fit, a smaller value of either information criteria suggests a better model. The information criteria penalize the fit with the number of parameter used.

The estimated parameters of the log likelihood ratio are not straightforward to interpret. This is true in particular for the heteroskedastic models where there are many higher order and interaction terms. We argue that a more constructive way to understand the results is to produce survival functions for distinct price levels which illustrate the behaviour of the model to distinct histories of prices. The original parameter estimates can be found in the Appendix in Tables 3 and 4.

4.2. Duration analysis

The probabilities for a household to switch from a variable-price contract to a fixed-price contract exactly at time τ, or after time τ, given a history of information $\mathcal{X}_\tau$, are

$$\Pr[Exit = \tau | \mathcal{X}_\tau, z_0] = \prod_{k=1}^{\tau-1} P_k^1 \cdot P_\tau^0, \tag{29}$$

$$\Pr[Exit > \tau | \mathcal{X}_\tau, z_0] = \prod_{k=1}^{\tau} P_k^1, \tag{30}$$

with $P_t^1 \equiv \Pr[D_t = 1 | \mathcal{X}_t, t, z_0]$ and $P_t^0 \equiv 1 - P_t^1$.

The last two expressions show the relationship between the conditional exit probabilities we discuss in Equation (8) and the probability that an exit from the variable-price contract takes place after or exactly at some time τ. Current, future and past probabilities to leave a variable-price contract are related, and we must have:

$$\Pr[Exit > \tau | \mathcal{X}_\tau, z_0] + \Pr[Exit = \tau | \mathcal{X}_\tau, z_0] = \Pr[Exit > \tau - 1 | \mathcal{X}_{\tau-1}, z_0], \tag{31}$$

since

$$\Pr[D_\tau = 1 | \mathcal{X}_\tau, \tau, z_0] + \Pr[D_\tau = 0 | \mathcal{X}_\tau, \tau, z_0] = 1,$$

That is, at any time t either the household keeps the variable-price contract or exits it (to a fixed-price one).[13]

The discrete survival function given in Equation (30) describes the survival probability up to time t given a specific history $\mathcal{X}_t$ of the information. Denote $\mathcal{P}(\mathcal{X})$ the sequence of survival probabilities given a complete history of information $\mathcal{X}$ from $t = 1$ to $t = +\infty$, that is:

$$\begin{aligned}\mathcal{P}(\mathcal{X}) =& (\Pr[Exit > 1 | \mathcal{X}, z_0], \Pr[Exit > 2 | \mathcal{X}, z_0], ..., \Pr[Exit > +\infty | \mathcal{X}, z_0]) \\ =& \big(\Pr[Exit > 1 | \mathcal{X}_1, z_0], \Pr[Exit > 2 | \mathcal{X}_2, z_0], ..., \Pr[Exit > t | \mathcal{X}_t, z_0], ..., \Pr[Exit > +\infty | \mathcal{X}, z_0]),\end{aligned}$$

where the second expression stresses that the t^{th} probability of $\mathcal{P}(\mathcal{X})$ depends on the history up to time t only. Clearly each element of $\mathcal{P}(\mathcal{X})$ is a probability, and furthermore, Equation (31) implies that the sequence is (weakly) decreasing. Hence, we can consider the expected sequence of survival probabilities, $\mathcal{P}$, over all possible histories of the information:

$$\mathcal{P} \equiv \mathop{\mathbb{E}}_{\mathcal{X}}[\mathcal{P}(\mathcal{X})].$$

Direct examination of the elements of $\mathcal{P}(\mathcal{X})$ suggests that the t^{th} element of $\mathcal{P}$ can be expressed as:

$$\mathcal{P}_t = \mathop{\mathbb{E}}_{\mathcal{X}}[\Pr[Exit > t | \mathcal{X}_t, z_0]] = \mathop{\mathbb{E}}_{\mathcal{X}_t}[\Pr[Exit > t | \mathcal{X}_t, z_0]]. \tag{32}$$

We can therefore use the parameter estimates of our to construct an estimate of $\mathcal{P}$ provided that we describe the data generating process which creates the histories $\mathcal{X}$.

Figure 8 illustrates such a construction for the homoscedastic model which accounts only for the history of price differentials and where the log odds are determined according to the modified Bayesian updating procedure we describe in Equation (5). To obtain the figure, we estimate a simple time series model on the observed history of the (deseasonalised log of the) variable prices (see appendix for more details). In this case, the simplest model takes the form of an AR(1) model. We then calculate the standardised residuals. We set the fixed price to 50 öre/kWh and we draw from the residuals independently and with replacement an alternative history of the variable price

[13] In a more realistic version, the discussion should account for exits to neither type of contract, i.e., the firm could loose the contract altogether to a competitor

innovations which we use together with the estimated model to generate an alternative history of the variable price from July 2009 until June 2014. Hence, the variable price and fixed price together allow us to generate an alternative history of the price differential. Figure 8 shows the effect of several (fifteen) such possible alternative histories on the survival probabilities (the dotted lines), as well as the average over all possible histories of the survival probabilities (the continuous line).

A first thing to notice is that the shape of the survival functions in Figure 8 all resembles the shape of the Kaplan-Meier survival function in Figure 7. In particular, there is a rather sharp decrease in the survival probability for short durations, and the survival function thereafter becomes almost flat. However, it is noticeable that distinct histories of the price differential yield a considerable range of the proportion that survive with a variable-price contract after sixty months since it ranges from below 40% to more than 80%. It is the timing of large positive shocks to the price differential during a particular history which determines variation to the proportion remain with a variable price contract. If the large shocks arise early, the proportion of survivors is larger than if the large shocks arise later in the history.

In Figure 9 we repeat the same exercise for different fixed prices (i.e., draw 250 histories for the variable price process, and generate the price differential given the fixed price, and then average over the conditional survival probabilities). The fixed prices we observe in the raw data range between 38 and 69 öre/kWh. However, we observe in Figure 1 that at the start of our sample period (June 2010) the fixed price rise from about 50 to 65 öre/kWh (February 2011) and decreases continuously thereafter to reach 40 öre/kWh by the end of our sample. The latest entry in our sample takes place in January 2012 when the fixed price was back at 50 öre/kWh. Hence we let fixed prices vary between 50 and 65 in increments of 5 öre since the observed entries to flexible price contracts arise only over that range.

Again, we note that the shape of the survival functions resembles the shape of the Kaplan-Meier estimator in Figure 7, but that there is substantial variation across different histories of prices. Further, over that range of fixed prices, the survival probabilities are ordered: given the data generating process that generates the variable price, a lower fixed price generates a sequence of smaller survival probabilities. Furthermore, the model suggests that much of the effect on the survival probabilities of changing the fixed price takes place between 50 and 60 öre/kWh. There is an effect above 60 öre/kWh but it appears less pronounced.

We repeat the same procedure and we generate the average survival sequence when the history of information contains the price differential and the temperature deviations from the household local monthly average. Hence, we estimate a time series model of the (log of the) variable price which conditions on last month variable price as well as on the current and the previous month deviation from the local temperature average (for the calculation presented this model is estimated assuming the household is located in Umeå). From the standardised residuals of this model, keeping the history of the local weather to its observed values, we generate alternative histories of the variable price. Together with each distinct value of the fixed price, we can then use the more elaborate model of switching to a fixed-price contract (which accounts for the history of the price differential and the history of temperature with the model of modified Bayesian updating) to generate the probabilities of survival with a variable-price contract.

Figure 10 shows the result of these calculations for an average household, i.e., a household at the sample mean for $z_0\hat{\pi}_0$ experiencing the temperature history of Umeå. As in the earlier simulation experiment, the sequences of survival probabilities are broadly convex functions of the contract time. The regular non-convexities to the survival probabilities that we observe in all cases reflect the (yearly) seasonal shocks to temperature. Furthermore, the survival functions for the distinct fixed prices are ordered and, as expected, show more response when the fixed price is small (50 öre/kWh) than when it is close to the maximum observed in the sample (65 öre/kWh). Five years after contracting, less than seventy five percent of households still remain with a variable-price contract if the fixed price is set 50 öre/kWh, while more than ninety percent would still have a variable-price contract if the fixed price is set at 65 öre/kWh. The more elaborate specification used here therefore suggests that the switching process responds differently to changes of the fixed price than the switching process derived from the homoscedastic model relying only on the price information. Relative to Figure 9, Figure 10 shows that the survival rates are more responsive to a price increase at higher prices (55 or 60) than to a price increase at the price of 50.

5. Conclusions

In this paper, we explore how sensitive the timing of transition away from the variable-price contract is to current and past prices. As far as we are aware, this paper is the first to explore this question for the choice between electricity contracts. To that effect, we propose a model for time series of binary decisions which allow for the dependence on the history of prices, or more generally a history of information. The model is consistent with the conventional decision theoretical framework. Compared to a less elaborate model which would rely only on current prices, our model illustrates how a potentially long sequence of past information, (in our case, prices and temperature) summarised into a parametrised log-likelihood ratio, can determine the choice to alter contract choice at any moment. This capture in a simple fashion the intuition that a decision maker needs to be convinced by observable facts before she agrees to change her mind about her current contract. We show furthermore that maximum likelihood estimation of the parameters of the decision model is feasible.

We provide several novel insights: First, we illustrate that whether recent evidence (i.e., the current price) favours the state of the world when the variable-price contract is preferable to the

Figure 8 "Average" Survival vs. history specific survival function

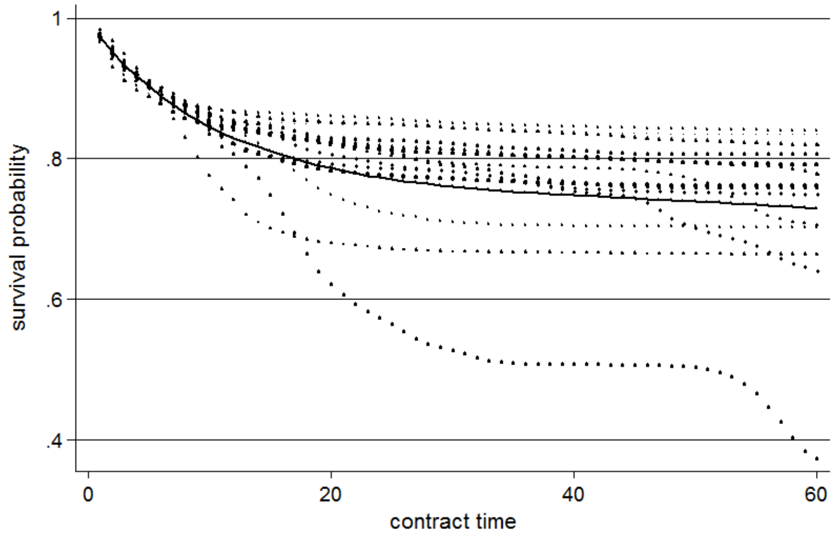

Note: The figure shows the "average" (over 250 possible histories of the variable price, continuous line) survival function (survival with a variable-price contract) at price 50 öre/kWh against the survival function for fifteen distinct variable price histories (dotted lines) based on the model specification which includes price information only and allows for information discounting.

Figure 9 "Average" Survival function, price information only, homoscedastic specification

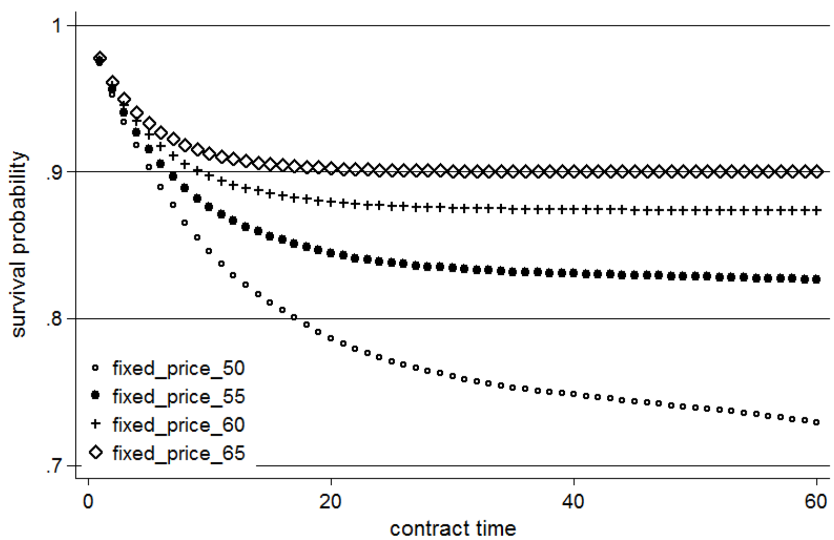

Note: The figure shows the "average" (over 250 possible histories of the variable price) survival function (survival with a variable-price contract) for different fixed prices based on the model specification which includes price information only and allows for information discounting.

Figure 10 "Average" survival function, price and temperature information, heteroscedastic information

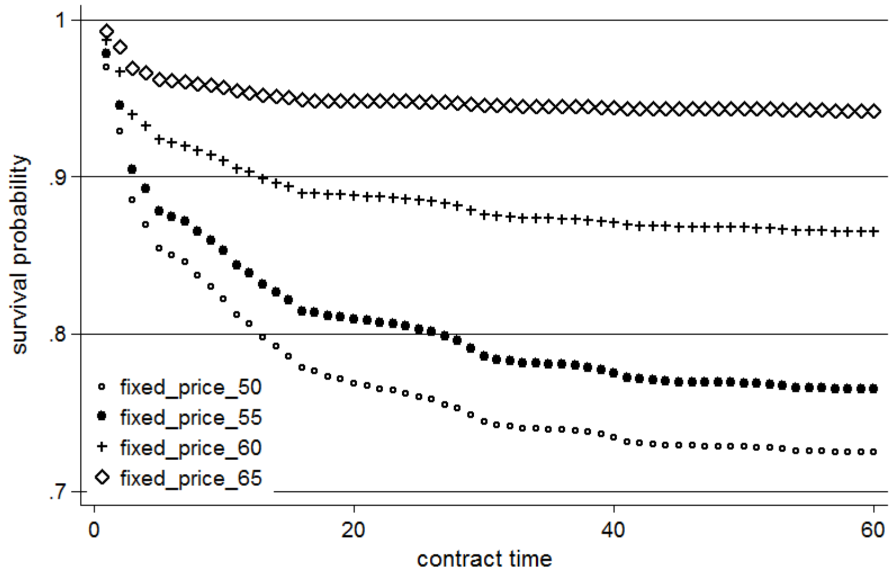

Note: The figure shows the "average" (over 250 possible histories of the variable price) survival function (survival with a variable-price contract) for different fixed prices based on the model specification which includes price and temperature information and allows for information discounting.

fixed-price contract (in the sense that it increases the posterior odds that $S = 1$) is not obvious, and that it depends crucially on the values of the autoregressive parameter in the price process. Therefore, it is not necessarily the case that households interpret higher variable prices as evidence of fixed-price contracts being more favourable. Our model suggests this to be the case irrespective of whether households believe that the autocorrelation coefficients and the variance of the innovation depends on the state of the world.

Second, we find that households do respond to new information, and that the seemingly small response to the price differential in the population may be a result of the characteristics of the process. In particular, we illustrate how alternative price histories may result in a substantial number of transitions. For example, in the homoskedastic specification, roughly seventy-five percent of households would still have a variable-price contract after sixty months if the fixed price is set at 50 öre/kWh, while ninety percent would still have a variable-price contract if the fixed price is set at 65 öre/kWh. The heteroskedastic specification (when the autocorrelation coefficients and the variance of the innovation depends on the state of the world) produces an even larger range of the proportion of households that remain on a variable-price contract after sixty months. The results also illustrate that most of the response to prices occur after a relatively short time, and that after the survival function is almost flat for longer durations.

From the retailers perspective, the results suggest that the retailer may modify the price differential (by changing the level of the fixed-price) to increase or decrease the switching rate, if that is beneficial to the retailer. However, such pricing strategies may have effects on customer churn that obviously need to be accounted for.

From a policy perspective, the finding that households do respond to prices by switching between contracts suggest that switching between contracts is an important margin of price responsiveness which should be accounted for when evaluating demand response for residential electricity demand. Although the switching rate appears to be small in the population, it may be substantial for alternative price processes, as already noted. In particular, the ongoing change in the production mix of electricity generation in the Nordic market, with an increasing share of intermittent generation, will most likely result in alternative price processes with more volatile prices, and possibly also changes in the price levels. The response to such alternative price processes may have substantial effects on the composition of share of households on each type of contract, which can have important policy implications: For example, if households to a larger extent switch to fixed-price contracts, this will reduce the overall response to short-run variation in the availability of electricity because these households then face no price-per-kWh variation.

The results in this paper, while novel and plausible, suffer from a few data-related drawbacks. In particular, the limited geographical variation in the sample calls for some caution when extrapolating the results to the entire Swedish population. The lack of variation at the household level for some socio-economic variables, and in particular temperature which, our study suggests, contains important information, is an obvious drawback of the data. On the other hand, many of the household characteristics are typically fixed in the short run, and should therefore have little effect on the the timing of transitions between electricity contract.

References

Benartzi S, Thaler RH (2007) Heuristics and biases in retirement savings behavior. *The Journal of Economic Perspectives* 21(3):81–104.

Borenstein S (2012) The redistributional impact of nonlinear electricity pricing. *American Economic Journal: Economic Policy* 56–90.

Brännlund R, Ghalwash T, Nordström J (2007) Increased energy efficiency and the rebound effect: Effects on consumption and emissions. *Energy Economics* 29(1):1–17.

Brännlund R, Karimu A, Söderholm P, et al. (2012) Elmarknaden och elprisets utveckling före och efter avregleringen: ekonometriska analyser. *CERE working paper* .

Braun M, Schweidel DA (2011) Modeling customer lifetimes with multiple causes of churn. *Marketing Science* 30(5):881–902.

Brennan TJ (2007) Consumer preference not to choose: Methodological and policy implications. *Energy Policy* 35(3):1616–1627.

Campbell JY, Cocco JF (2003) Household risk management and optimal mortgage choice. Technical report, National Bureau of Economic Research.

Ching AT, Erdem T, Keane MP (2013) Invited paper on learning models: An assessment of progress, challenges, and new developments. *Marketing Science* 32(6):913–938.

Damsgaard N, Green R, Johansson B (2005) Den nya elmarknaden: Framgång eller misslyckande? *SNS förlag* .

DeGroot MH (2005) *Optimal Statistical Decisions* (John Wiley & Sons, Inc.), ISBN 9780471729006, URL `http://dx.doi.org/10.1002/0471729000`.

Dhillon US, Shilling JD, Sirmans C (1987) Choosing between fixed and adjustable rate mortgages: Note. *Journal of Money, Credit and Banking* 19(2):260–267.

Ek K, Söderholm P (2008) Households' switching behavior between electricity suppliers in Sweden. *Utilities Policy* 16(4):254–261.

Ericson T (2011) Households self-selection of dynamic electricity tariffs. *Applied Energy* 88(7):2541–2547.

Fowlie M, Greenstone M, Wolfram C (2015) Are the non-monetary costs of energy efficiency investments large? Understanding low take-up of a free energy efficiency program. *The American Economic Review* 105(5):201–204.

Goett AA, Hudson K, Train KE (2000) Customers' choice among retail energy suppliers: The willingness-to-pay for service attributes. *The Energy Journal* 21(4):1–28.

Goettler RL, Clay K (2011) Tariff choice with consumer learning and switching costs. *Journal of Marketing research* 48(4):633–652.

Gupta S, Hanssens D, Hardie B, Kahn W, Kumar V, Lin N, Ravishanker N, Sriram S (2006) Modeling customer lifetime value. *Journal of Service Research* 9(2):139–155.

Ito K (2014) Do consumers respond to marginal or average price? Evidence from nonlinear electricity pricing. *The American Economic Review* 104(2):537–563.

Juliusson EA, Gamble A, Garling T (2007) Loss aversion and price volatility as determinants of attitude towards and preference for variable price in the Swedish electricity market. *Energy Policy* 35(11):5953–5957.

Klemperer P (1987) Markets with consumer switching costs. *The Quarterly Journal of Economics* 102(2):375–394.

Krishnamurthy CKB, Kriström B (2015) A cross-country analysis of residential electricity demand in 11 OECD-countries. *Resource and Energy Economics* 39:68–88.

Lambrecht A, Seim K, Skiera B (2007) Does uncertainty matter? Consumer behavior under three-part tariffs. *Marketing Science* 26(5):698–710.

Lambrecht A, Skiera B (2006) Paying too much and being happy about it: Existence, causes, and consequences of tariff-choice biases. *Journal of Marketing Research* 43(2):212–223.

Littlechild S (2006) Competition and contracts in the Nordic residential electricity markets. *Utilities Policy* 14(3):135–147.

Miravete EJ (2002) Estimating demand for local telephone service with asymmetric information and optional calling plans. *The Review of Economic Studies* 69(4):943–971.

Miravete EJ (2003) Choosing the wrong calling plan? Ignorance and learning. *The American Economic Review* 93(1):297–310.

Revelt D, Train K (2000) Customer-specific taste parameters and mixed logit: Households' choice of electricity supplier. *Department of Economics, UCB, working paper* .

SCB (2014) Renegotiations and switching of electricity contracts (Omförhandling och byten av elavtal). Technical report, Statistics Sweden.

Vesterberg M, Kiran B Krishnamurthy C (2016) Residential end use electricity demand: The implications for real time pricing in sweden. *Energy Journal* 37(4):141–164.

Vigna SD, Malmendier U (2006) Paying not to go to the gym. *The American Economic Review* 96(3):694–719.

Wilson CM, Price CW (2010) Do consumers switch to the best supplier? *Oxford Economic Papers* gpq006.

Yang Y (2014) Understanding household switching behavior in the retail electricity market. *Energy Policy* 69:406–414.

6. Appendix

Table 3 **Parameter estimates for specifications with univariate homoskedastic information (spec 1, 2 and 3) and bivariate homoskedastic information (spec 3, 4 and 5)**

	(1)	(2)	(3)	(4)	(5)	(6)
$x_{2,t}$	0.123	1.045***	1.297***	-1.406***	-0.407	-0.178
	(0.169)	(0.177)	(0.208)	(0.188)	(0.215)	(0.239)
$x_{2,t-1}$	-0.960***	-1.566***	-1.864***	-0.670***	-1.133***	-1.546***
	(0.169)	(0.176)	(0.213)	(0.171)	(0.181)	(0.216)
$x_{1,t}$				-0.153***	-0.106***	-0.118***
				(0.008)	(0.009)	(0.010)
Constant	4.719***	3.622***	3.690***	4.367***	3.854***	3.819***
	(0.371)	(0.373)	(0.373)	(0.374)	(0.376)	(0.376)
Trend		0.126***	0.152***		0.0708***	0.107***
		(0.007)	(0.013)		(0.008)	(0.013)
Discount			0.980**			0.978***
			(0.007)			(0.005)
Number obs.	83786	83786	83786	83786	83786	83786
Log-likelihood	-3857.7	-3624.4	-3620.3	-3595.5	-3548.9	-3538.9
AIC	7765.3	7300.7	7294.7	7243.0	7151.8	7133.9
BIC	7998.7	7543.4	7546.8	7485.7	7403.9	7395.3

Standard errors in parentheses
* $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

Table 4 Parameter estimates for specifications with heteroskedastic information

	(7)	(8)	(9)
$x_{2,t}^2$	0.609	10.55***	19.29***
	(1.705)	(2.404)	(3.632)
$x_{2,t}x_{1,t}$	-0.234	0.334	1.057**
	(0.190)	(0.218)	(0.343)
$x_{2,t}x_{1,t-1}$	-0.580**	-0.676***	0.010
	(0.192)	(0.197)	(0.258)
$x_{2,t}x_{2,t-1}$	-12.76***	-18.13***	-21.02***
	(2.673)	(2.910)	(3.507)
$x_{2,t}$	0.684*	0.351	1.512***
	(0.304)	(0.321)	(0.414)
$x_{2,t-1}^2$	4.010**	9.945***	14.90***
	(1.542)	(1.827)	(2.332)
$x_{2,t-1}$	-2.331***	-0.327	-0.944*
	(0.296)	(0.386)	(0.455)
$x_{1,t}^2$	0.016*	0.029***	0.065***
	(0.007)	(0.007)	(0.012)
$x_{1,t}x_{1,t-1}$	-0.053***	-0.065***	-0.049**
	(0.011)	(0.01)	(0.015)
$x_{1,t}x_{2,t-1}$	0.488***	0.329**	1.059***
	(0.122)	(0.125)	(0.189)
$x_{1,t}$	-0.076***	-0.055*	-0.094***
	(0.021)	(0.022)	(0.027)
$x_{1,t-1}^2$	-0.014	0.002	0.0379**
	(0.008)	(0.008)	(0.012)
$x_{1,t-1}x_{2,t-1}$	-0.809***	-0.294	-0.093
	(0.171)	(0.197)	(0.265)
$x_{1,t-1}$	0.021	0.092***	0.174***
	(0.022)	(0.024)	(0.032)
Constant	3.755***	3.899***	3.296***
	(0.378)	(0.378)	(0.382)
Trend		-0.331***	-0.467***
		(0.040)	(0.051)
Discount			0.894***
			(0.015)
Number obs.	83786	83786	83786
Log-likelihood	-3396.8	-3360.1	-3323.9
AIC	6867.6	6796.2	6725.8
BIC	7213.0	7151.0	7089.9

Standard errors in parentheses
* $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

Table 5 Autoregressive models, (ln) variable price

	Price only	Price and temperature
Constant	0.060	0.044
	(0.095)	(0.079)
ln Variable price (t-1)	0.786***	0.769***
	(0.089)	(0.096)
Temperature (t)	-	-0.030***
		(0.007)
Temperature (t-1)	-	-0.018***
		(0.007)
σ	0.156***	0.136***
	(0.013)	(0.011)
Log-likelihood	26.37	34.79
Q-test stat	4.36	4.95
p-value (10 d-o-f)	0.93	0.90

Standard errors in parentheses
* $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

Note: i) The table report maximum likelihood estimates of the AR(1) models used to construct the alternative histories of the variable price. The models are estimated over the sample period of interest from June 2009 to June 2014 (61 observations).

ii) The Q-test statistics tests the null hypothesis that the residuals is the realisation of a white noise process. We test the first 10 autocorrelations.