Electricity price predictability and quantifying customer suitability

Problem submitted by: Paddy Finn, Crystal Energy.

Report prepared by:

Mel Devine[1,*], James Gleeson[1].

Workshop participants: Rob Arculus[2], Mark Christiansen[3], Mel Devine[1], Conor Dempsey[1], James Fennell[1], James Gleeson[1], Geoff Hunter[4], Kevin O′Sullivan[1], Ben Nuttall[5].

[1] MACSI, Department of Mathematics and Statistics, University of Limerick, Ireland.

[2] London School of Economics, London, UK.

[3] Trinity College Dublin, Dublin, Ireland.

[4] Limerick Institute of Technology, Limerick, Ireland.

[5] Manchester Metropolitan University, Manchester, UK.

September 14, 2011

*Author for correspondence: Mel Devine, MACSI, Department of Mathematics and Statistics, University of Limerick, Limerick, Ireland. mel.devine@ul.ie

Abstract

This report summarises progress made towards the problems submitted by Crystal Energy at the 82^{nd} European Study Group with Industry. Crystal Energy purchases power from the Irish Single Electricity Market, and supplies it to small and medium sized business customers at half-hourly variable rates. They posed two problems to the study group. The first was to analyse the predictability of these electricity prices while for the second problem the group were asked to develop a metric that would help identify customers that are suitable for these flexible tariffs. The group's analysis found that large price spikes and large errors between actual and predicted prices were more likely to occur at certain times. Previously, Crystal Energy supplied their customers with price predictions only. Using the results from the study group, Crystal Energy will now be able to provide their customers with information on the errors associated with price predictions as well as information on the likelihood of price spikes at a given time. In relation to the second problem, the group successfully developed a measure that is able to quantify the suitability of different customers to the flexible tariffs offered by Crystal Energy.

1 Introduction

Crystal Energy is a licensed supplier of electricity from the Irish Single Electricity Market to business customers. Crystal Energy's customers are charged a rate proportional to the wholesale price (the "system marginal price" (SMP)) of the electricity supplied; this is typically lower at low-demand times of the day, and higher at times of peak usage. In contrast, most small and medium-sized business customers would receive only a flat-rate pricing structure from the dominant national electricity suppliers. Thus Crystal Energy encourages efficient demand-side management (scheduling of electricity demand at times when electricity prices are cheaper) by its customers, and their ideal customers are those who have the ability to reschedule their demand (or "load") away from high-demands times and into low-price periods.

The Single Electricity Market Operator (SEMO) supplies market information, in the form of a day-ahead predicted ("ex-ante") price, to Crystal Energy. This ex-ante price is determined by the solution of a mixed integer programming problem that gives the optimal scheduling of available generators (the "unit-commitment problem"). Roughly speaking, the base price of electricity (the "shadow price") in each half-hourly time interval is set by the most expensive generator that is required to run to meet demand—the most costly generators are scheduled to run only at times of peak demand, causing the higher prices at these times. However, the details of the unit-commitment optimization constraints are more complicated than this, involving ramp-up times for generators, distinctions between cold and hot starts, special treatment for pumped storage units, etc. [1, 2]. Moreover, the SEM Trading and Settlement Code requires that all generators recoup their running costs, leading to an addition to the shadow price called "uplift". The values of the uplift can be very large and concentrated on small intervals of time, and this leads to a price time-series that exhibits infrequent but large-amplitude "spikes". The uplift process was examined in a previous Study Group [3, 4].

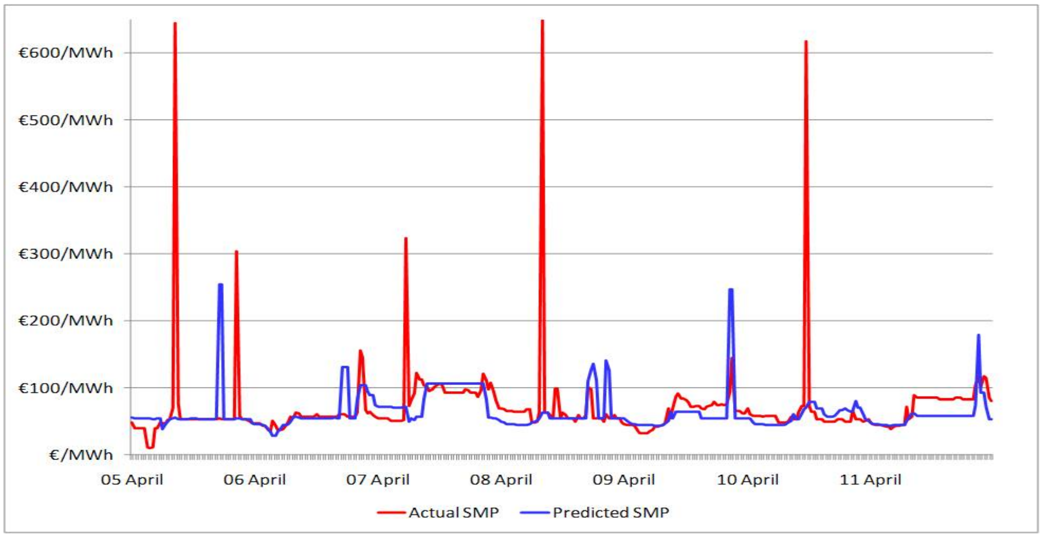

Figure 1: Actual and predicted SMP

The ex-ante price for the day ahead is published at 8pm each evening, and Crystal Energy pass on this information to their customers. A typical prediction is shown in Figure 9: the ex-ante price (in Euros per megawatt hour) is plotted for each half-hourly time interval from midnight to midnight. Customers may then use this information to reschedule their electricity load for the day ahead in order to avoid times of high predicted price. However, the ex-ante price and the actual ("ex-post") price are not always equal: see Figure 1 for an example of a week where substantial differences occurred. The errors in prediction may be due to one of a number of factors, such as errors in the demand prediction, errors in wind forecast, or unexpected outages of generators. The differences between predicted and actual prices pose a substantial risk to Crystal Energy's customers, as it is possible that they could reschedule load to period where price is expected to be low, only to have an unanticipated price spike occur there, leading to an unanticipated increase in their electricity cost.

The first problem posed by Crystal Energy to the Study Group was to analyze the predictability of the price. In particular, Crystal Energy sought a method for quantifying the risk of higher-than-predicted prices, and a mechanism for presenting this information in a readily understandable format. As a second problem, the Study Group was asked to design a metric for identifying customers whose historical demand profile would render them particularly suitable for Crystal Energy's tariff. The metric should identify those customers whose electricity usage is concentrated in the lower-price time intervals: these customers would benefit most from the SMP-following price offered by Crystal Energy.

The remainder of this report is structured as follows. In Section 2 we describe the statistical analysis of the price-prediction errors, and of the large spikes which cause these errors. Section 3 presents the application of the statistics from Sec. 2 to provide information to Crystal Energy's customers, thus solving the first problem. In Section 4 we explain the metric developed for the second problem, the quantification of customer suitability.

1.1 Data

Prices within the Single Electricity Market, both predicted and actual, are given on half-hourly time intervals. The data provided by Crystal Energy included the predicted SMP and actual SMP for each half-hourly time period over the three and half years from 01/01/2008 to 01/06/2011.

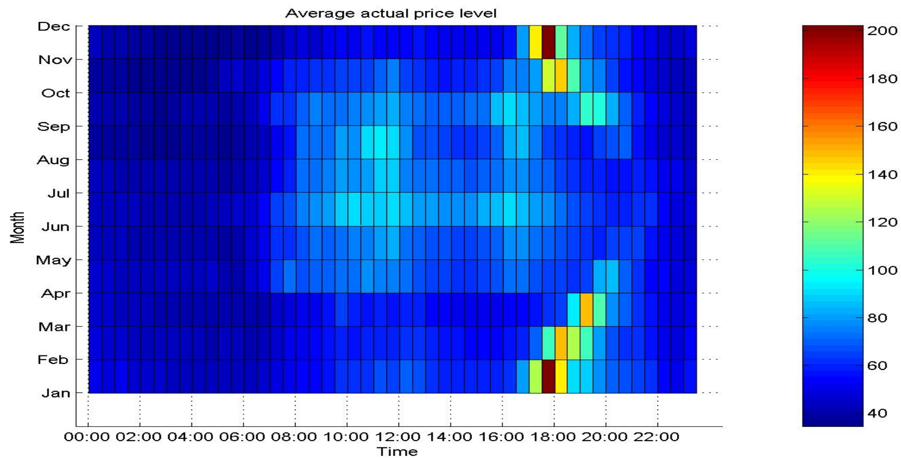

Figure 2: Heat map of the average actual price for each month and for each half hour time period in the day

Following discussion with Crystal Energy, the group decided that the data analysis should respect both the time-of-day timescale, and the time-of-year timescale, as both aspects of seasonality are important for electricity prices. Figure 2 shows a heat map of the actual SMP, where each element in the 12×48 grid represents the average price for a given month and a given half-hour time slot. The number of data points contributing to each grid value thus varies from approximately 90 (three years, with about 30 days per month) to 120 (for January through June, where four years of data were available). A distinct seasonal pattern is observable in Figure 2; this is analysed further below.

2 Analysis

In this section we present two related but distinct approaches taken by the group to analyse the price data sets. Section 2.1 shows results for the magnitude of the difference between the predicted and actual SMP values (the "error"), while in

Section 2.2 the accuracy in predicting price spikes is examined.

2.1 Error analysis

To quantify the difference between the predicted and actual values of the SMP for a given half-hour time period, we define the error for that time period to be

$$error = SMP_{predicted} - SMP_{actual}. \tag{1}$$

Figure 3 shows a heat map of the average absolute values of this error measure. Note the striking resemblance to the seasonal pattern observed in Figure 2: this indicates that high error values occur at similar times of the day to high price values. This motivates us to further examine the mechanism for large errors (and large prices) in the spike analysis of Section 2.2 below.

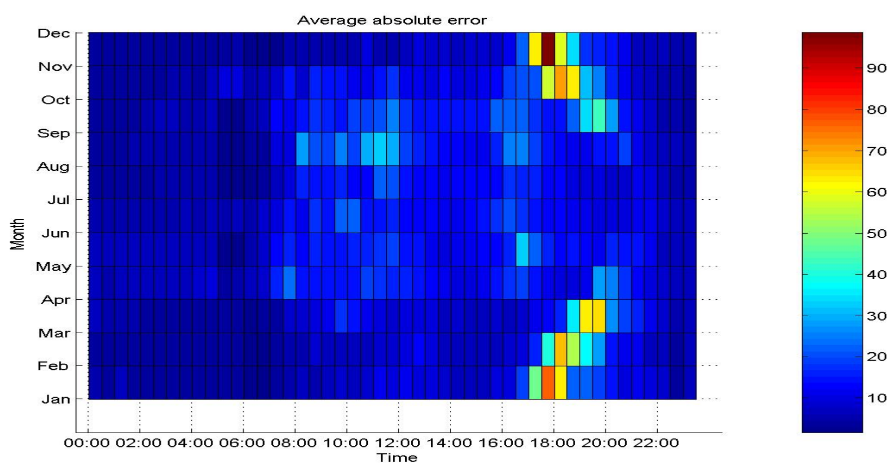

Figure 3: Heat map of the average absolute errors for each month and for each half hour time period in the day

A histogram of the error values (for all months and all time periods) is shown in Figure 4, and moments of the distribution are given in Table 1. Although the distribution of error values is approximately symmetric about zero, it has much fatter tails than a normal distribution: analysis of the tails (not shown) indicates that they have a power-law (algebraic) decay for large error values.

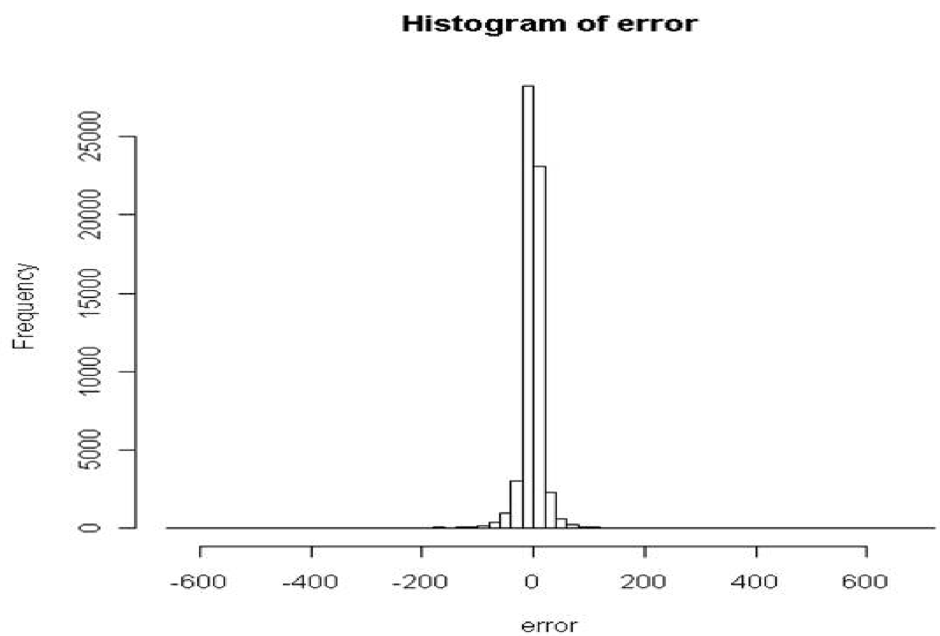

Figure 4: Histograms of the errors

Mean	-1.29
Median	0.00
Standard deviation	28.32
Minimum	-647.00
Maximum	720.00
Skewness	-1.71
Kurtosis	90.20

Table 1: Statistics of the error values

2.2 Price spike prediction

As mentioned in Section 1, Crystal Energy encourages their customers to move their load away from times of high prices. In particular, they encourage them to move away from large price spikes. In Figure 1 we can see an example of four large price spikes in the actual SMP data. In this section, we detail the analysis the group undertook on large price spikes and the analysis of how good the predicted SMP are at anticipating them. After considering a number of different options, the group defined a price spike as being above two times the average price of the previous 336 timesteps (i.e. two weeks).

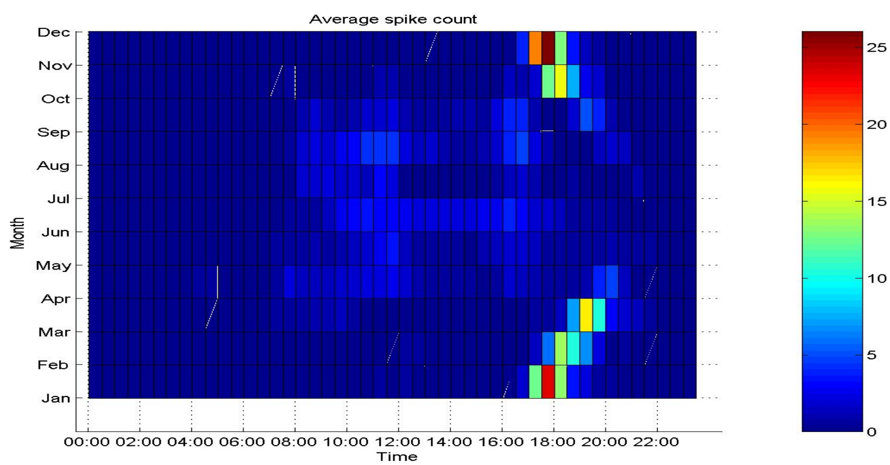

Figure 5: Heat map of the average number of spikes for each month and for each half hour time period in the day

Figure 5 shows the average number of price spikes using this metric for each half-hour period in the day, for each month. From this graph we can see that there is, on average, very few price spikes from 22:00 in the evening to 08:00 the following morning. The largest number of price spikes, on average, occur in the winter time at around 17:00 or 18:00 (i.e. dusk). The large number of price spikes remains for the spring and autumn months, however they occur at later times in the evening. This corresponds with the change in time of dusk for these months.

In the summer months, there are on average very few price spikes in the evening. Large price spikes tend to occur between morning and lunch time during this time. The locations of the large price spikes in Figure 5, are clearly close to the positions with high prices and large errors seen in the heat maps in Figures 2 and 3 respectively.

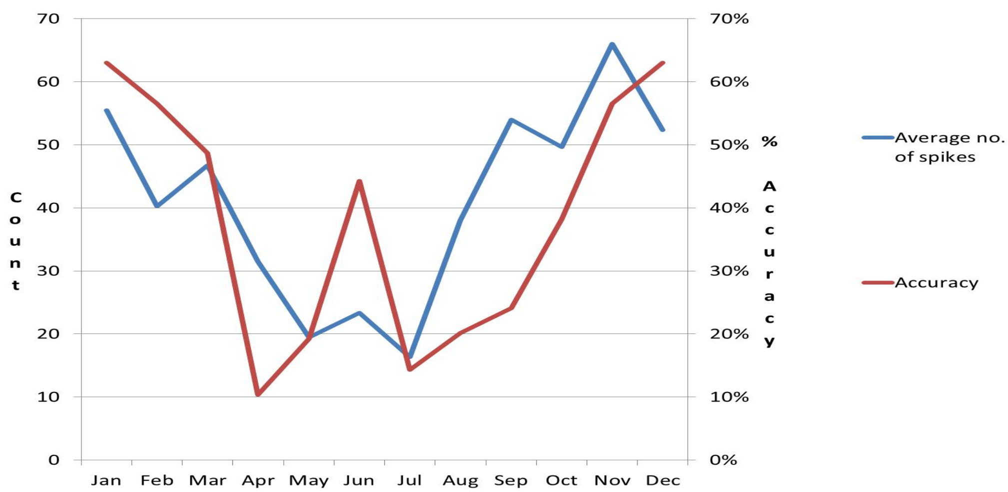

Figure 6: Average number of observed spikes per month plus the % accuracy in correctly predicting them.

Figure 6 shows the average number of observed spikes per month (blue line) as well as the the percentage accuracy in correctly predicting these spikes per month (red line). A spike was classified as correctly predicted if it occurred in the same time period in both the actual and predicted SMP data. The sizes of the spikes were not considered as long as they were both above the threshold defined above. This graph shows a seasonal effect to the average number spikes per month. On average, more spikes occur in the winter compared with the summer time. Figure

6 also shows that the percentage accuracy in correctly predicting spikes also has a seasonal dependence. The percentage accuracy in prediction tends to be higher in the winter months than in the summer months. In both January and December, the number of spikes that were correctly predicted was approximately 63%. By April this accuracy dropped to approximately 10%. Despite increasing to 44% in June the percentage accuracy does not go over 50% again until November. As mentioned in the introduction, Crystal Energy each evening supply their customers with predicted SMPs for the following day. If the customer sees a predicted price spike they will attempt, if possible, to move their load away from that time in order to significantly reduce costs. However, this analysis shows that there is a large percentage of price spikes that are not anticipated by the predicted SMP. Thus, when the customer is moving their load, they may be moving it to a time when a price spike may actually occur.

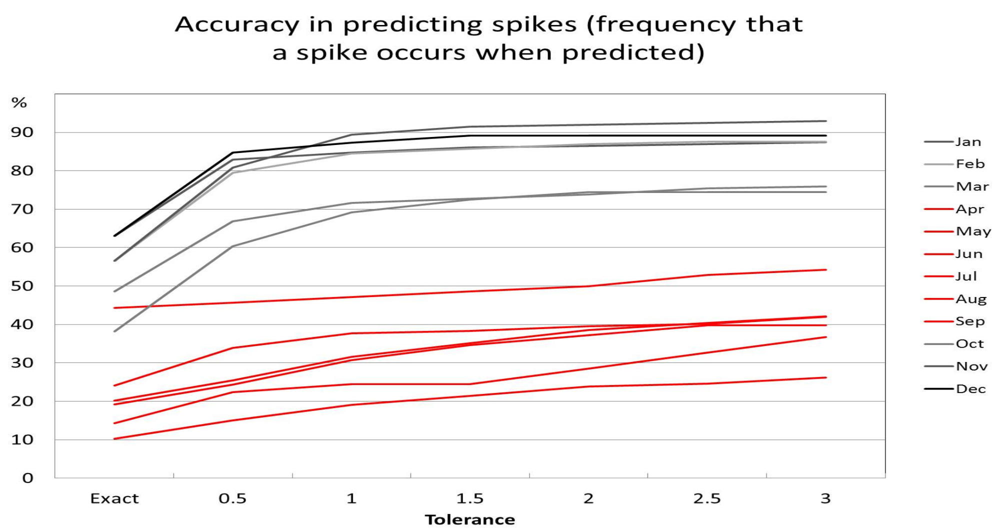

Figure 7: Accuracy of correctly predicting price spikes by month.

Figure 7 shows the accuracy in predicting price spikes, as a function of the ”tolerance” time interval. A zero tolerance means that if a spike was predicted at, for example 16:00, then for it to be classified as correctly predicted, the actual spike must also have occurred at 16:00 on the same day. A tolerance of one hour means however, that if a spike was predicted at a certain time but it actually occurred within an hour either side of that time, then it would also be classified as correctly predicted. For example, if a spike was predicted for 18:00 but it actually occurred between 17:00 and 19:00, then it would still be classified as correctly predicted.

As we've seen with Figure 6, Figure 7 clearly shows that spike prediction accuracy tends to be higher in the winter compared with the summer. (The summer months are in red while the winter months are in black.) It also shows that when the error tolerance is increased by a half an hour there is a significant improvement in the percentage accuracy across each month. Further improvements can be seen when the error tolerance is increased to one hour. For example, in January the percentage accuracy improves from 63% to 85% with a one hour tolerance. When the error tolerance is increased to one and half hours there are only minor improvements in the percentage accuracy. The same can be said when the tolerance is increased to two, two and half, and three hours. Figure 7 also shows that these improvements are larger in the winter months than in the summer months. For example, with a one hour tolerance, the percentage accuracy improves from an average of 49% to 79% in the winter months while the corresponding improvement in the summer months is from 16% to 31%.

This information can help Crystal Energy's customers' decision making process. Following from this analysis, they now know that if a spike is predicted for a given time, it may not occur at that time, but it may occur within a certain period around that time with a certain probability depending on the time of year. Previously, if a customer saw a price spike predicted at a given time, they may have moved their large load an hour or two earlier or later. However, this analysis shows there is still a high probability that the large price spike may occur around the predicted time. This is particularly true in winter.

3 Application of this analysis

The results of Section 2 were used by the group to provide extra information to Crystal Energy's customers in an accessible and user-friendly interface. Previously, Crystal Energy supplied its customers with a day-ahead prediction of the SMP: a typical screen seen by the customer is shown in Figure (8).

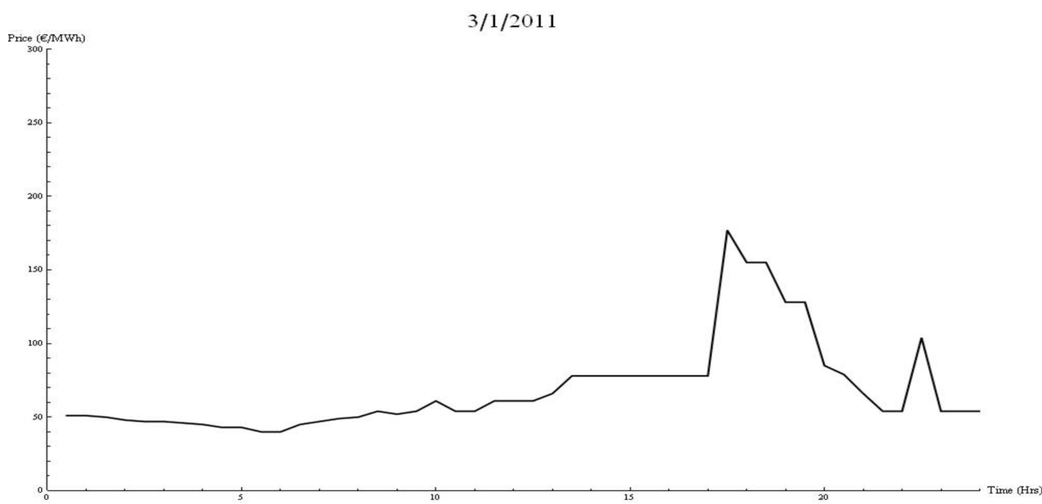

Figure 8: Predicted (day-ahead) price series for 3/1/2011, as currently supplied to Crystal Energy customers.

Using the analysis presented in Figures 3 and 5, we can now give an indication of the likely accuracy of the price prediction, from the empirical error distributions for the particular month and for each half-hourly time period. A first measure of accuracy is the mean positive error, and the mean negative error, for the time period. The dark green region in Figure 9 indicates prices which lie between the values defined by adding the mean positive error, and mean negative error, to the predicted SMP.

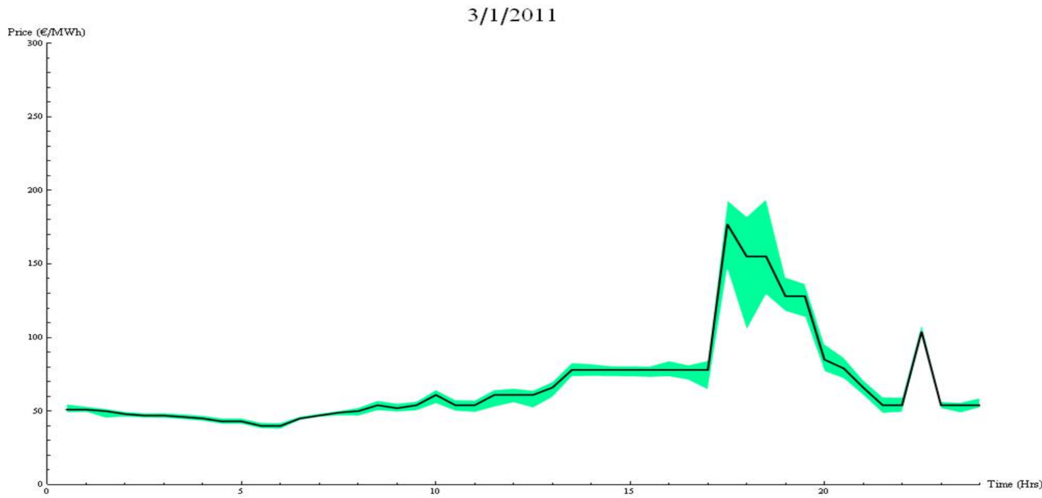

Figure 9: Predicted price series plus mean positive and negative errors for each time point for 3/1/2011. The dark green region indicates the uncertainty associated with the predicted price at that time.

This "uncertainty" region is widest at the peak time of the day, as this is the most likely times for spikes, and hence large errors, to occur. Since the distribution of errors is fat-tailed, it is important to give further information on the moments.

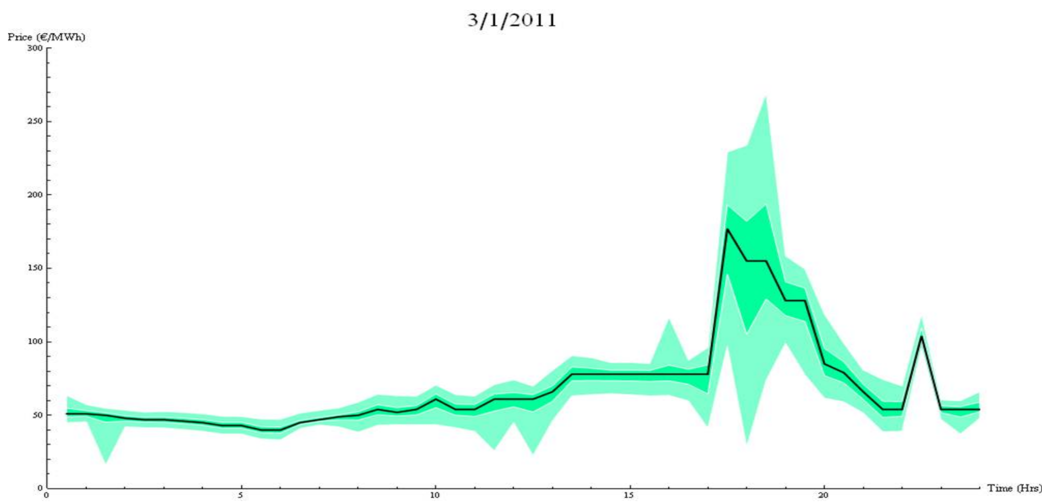

Figure 10: As Fig. 9, but including standard deviations for positive and negative errors (light green regions).

This is provided in Figure 10, where the light green region adds one standard deviation (positive or negative) error to the dark green region of Figure 9. Note that the light green region is of larger extent, and so gives a more visible indicator of time periods where errors are more or less likely.

Figure 11 adds the actual SMP time series to the day-ahead predictions of Figure 10.

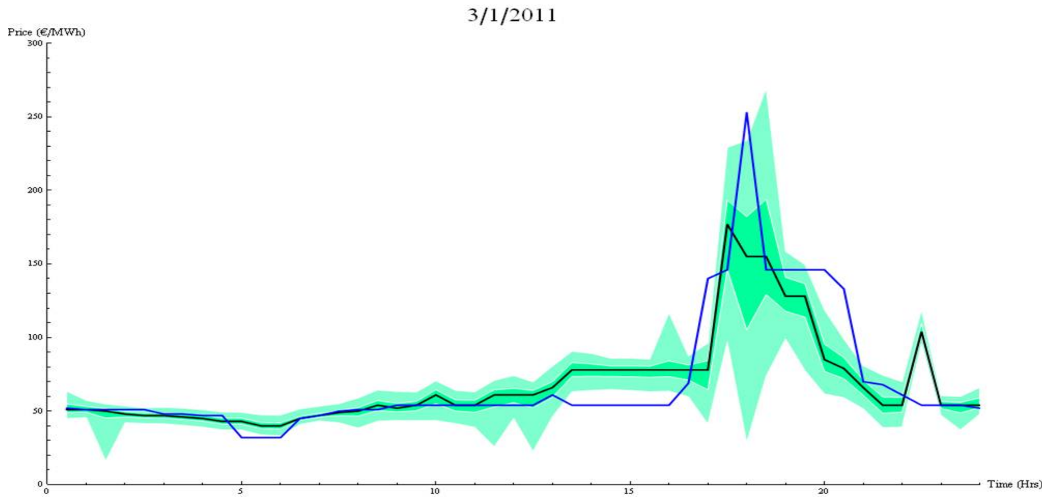

Figure 11: As Fig. 10, but showing the actual SMP for 3/1/2011.

Note the significant difference between the predicted and actual prices at 18:00, with relatively small errors at other times. A customer who had the benefit of the enhanced day-ahead analysis of Figure 10 could have anticipated that the actual SMP would behave in roughly this manner, and may have moved load to low-price times with relatively low risk.

As the analysis is based purely on historical data, there will of course be occasions when the actual SMP has large spikes which are not anticipated by our "uncertainty region": Figure 12 shows an example of such a day. However, overall Crystal Energy believe that this analysis provides valuable extra information to their customers, which will enable them to better implement demand-side management.

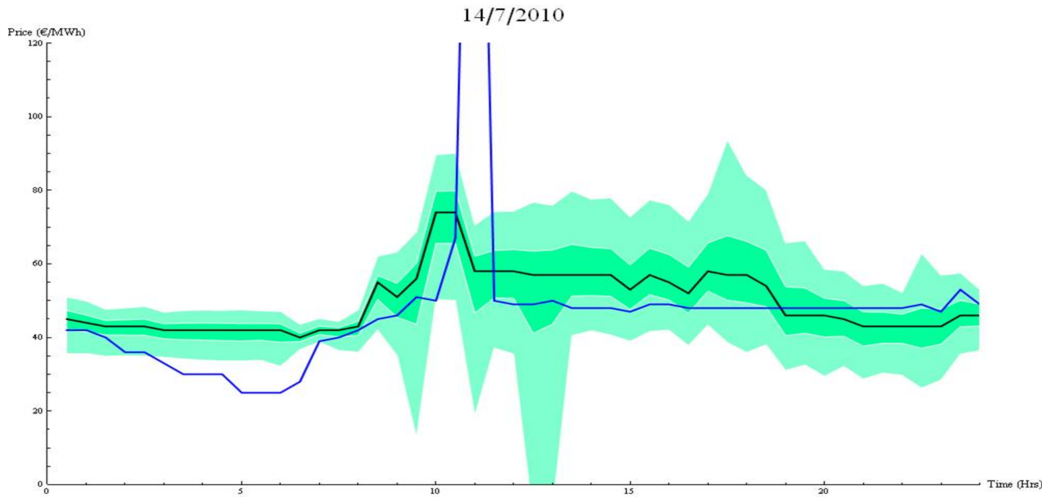

Figure 12: Predicted and actual prices, along with uncertainty regions, for 14/7/2010.

The spike-based analysis of Section 2.2 may also be used to further elaborate the day-ahead information available to the customer. Suggested presentation modes for spike probabilities within user-defined tolerances (i.e., time frames within which they can move load) are shown in Figures 13–15.

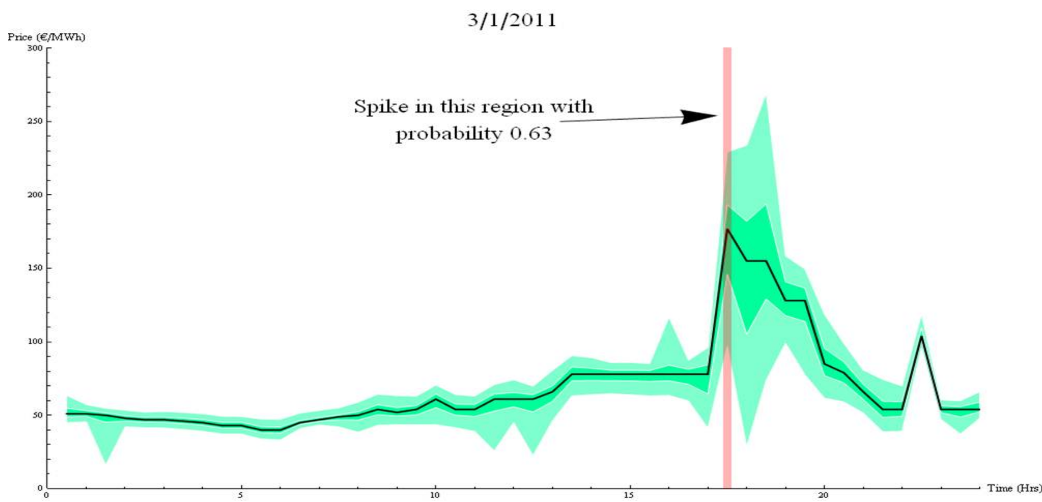

Figure 13: Predicted price series with error bounds along with probability of spike occurring during a half hour period.

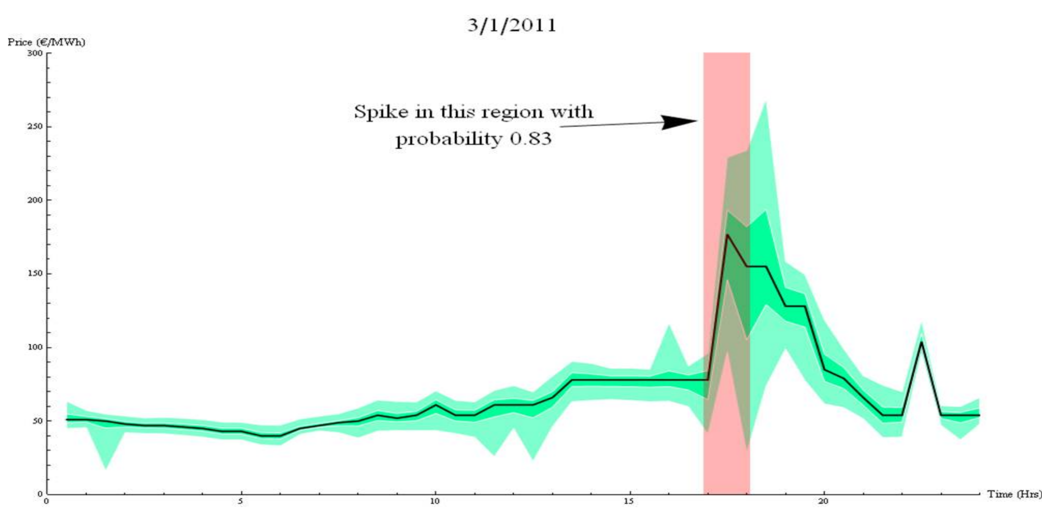

Figure 14: As Fig. 13, but with probability of spike occurring in one and an half hour period.

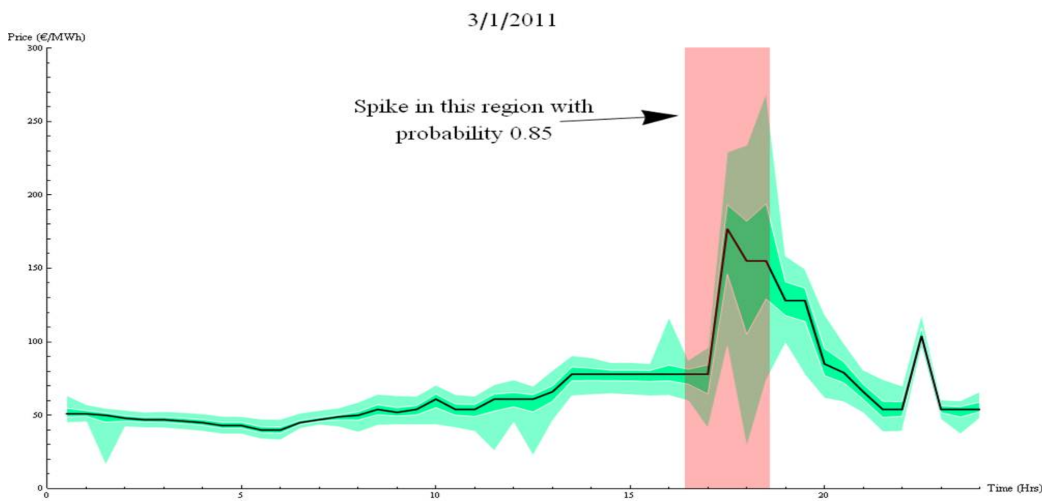

Figure 15: As Fig. 13, but with probability of spike occurring in two and an half hour period.

4 Customer suitability

As mentioned in the introduction, the second problem brought by Crystal Energy to ESGI 82 was to design a metric that helped Crystal Energy decide whether or not a customer was suitable for their flexible tariffs. Prior to the study group, Crystal Energy did not have any quantitative measure that allowed them to decide this. They relied on graphs such as figures (16) and (17) in order to establish how suitable a potential customer was.

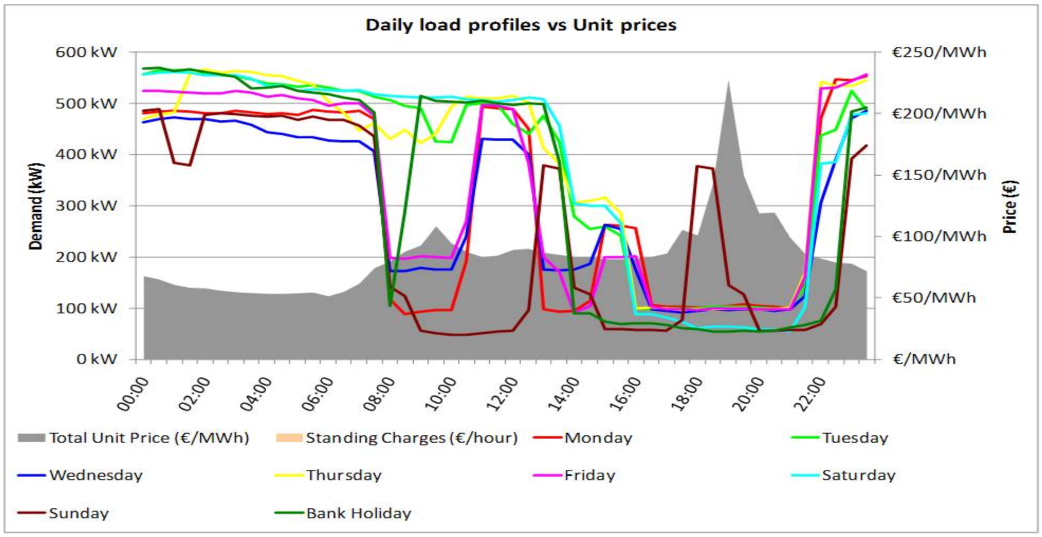

Figure 16: Load profile of a customer with a cold storage facility

For example, Figure 16 shows the average load profile for the seven days of the week at each half-hourly time period, for a customer of Crystal Energy. It also includes the average load for bank holidays as well as the average SMP (grey background) at each half-hourly time period. This customer operates a cold storage facility and from the graph we can see that the majority of its load comes in the morning to early-afternoon time, i.e., times of relatively low prices. When the price increases in the evening time, this cold storage facility customer has a relatively low load.

From this, we can easily conclude that this customer would be suitable for Crystal Energy's flexible tariffs as they could save a significant amount of money on their electricity bill. The analysis below will confirm this conclusion.

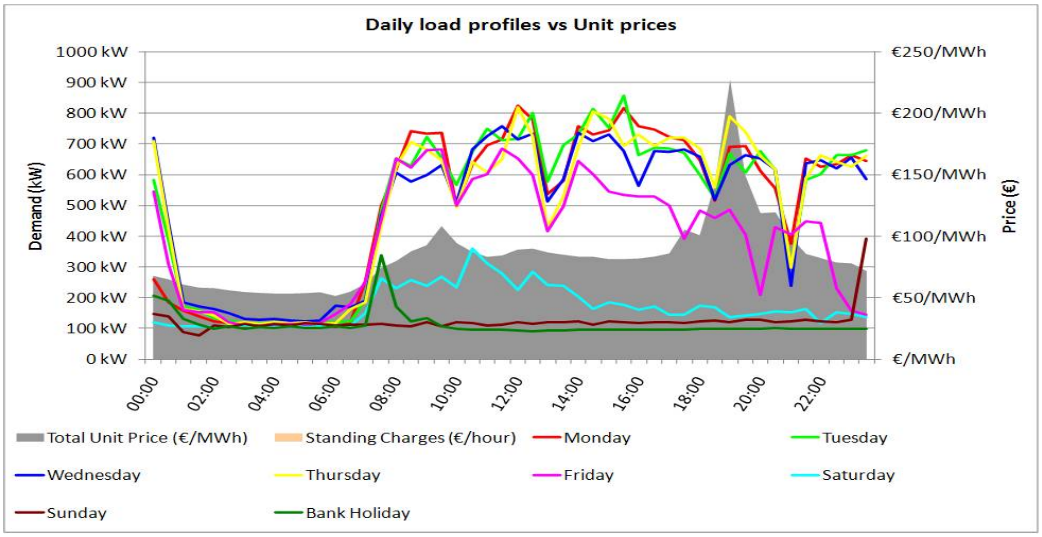

Figure 17: Load profile of some customer

However, it is not always possible to judge from these type of graphs whether a customer is suitable or not. Figure 17 shows the same type of graph for a different customer of Crystal Energy. This customer is a manufacturing company. The majority of their load occurs during the day. They have very little load during the night. Similar to the cold storage facility, this customer has a relatively large amount of load from the morning to afternoon time. In contrast however, the load remains relatively high for the high-priced evening time. Thus, it is difficult to conclude from this graph whether or not this customer is suitable for the flexible tariffs offered by Crystal Energy as they have relatively large loads at times of both high and low prices.

4.1 Load weighted average price

One quantitative measure suggested to the group, by Crystal Energy, was to calculate the Load Weighted Average Price (LWAP) for both the customer and the total system. The customer's LWAP was calculated using the following formula:

$$LWAP_{Customer} = \frac{\sum_t (Price_t \times CustomerLoad_t)}{\sum_t CustomerLoad_t}, \quad (2)$$

where $Price_t$ and $CustomerLoad_t$ represented the price and the load of the customer at time t. Similarly, the total system's LWAP was calculated using the following formula:

$$LWAP_{TotalSystem} = \frac{\sum_t (Price_t \times TotalSystemLoad_t)}{\sum_t TotalSystemLoad_t}, \quad (3)$$

where $TotalSystemLoad_t$ represented the total amount load on the system at time t. Once these two measures were calculated over the same time period they could be compared with each other. If the customer's LWAP was higher than the total system's (i.e. if $\frac{LWAP_{Customer}}{LWAP_{TotalSystem}} > 1$) this indicated that the customer tended to have more load at times of relatively high prices and was thus not suited to the flexible tariffs offered by Crystal Energy. Conversely, if the customer's LWAP was lower than the total system's (i.e. if $\frac{LWAP_{Customer}}{LWAP_{TotalSystem}} < 1$) this indicated that the customer tended to have more load at times of relatively lower prices, meaning they would be suited to the flexible tariffs.

To test this measure, the LWAP of four different customers were calculated and compared with the LWAP of the total system. Crystal Energy provided the data on the load (in kilowatts) used by four different customers at each half-hour time-step. The time period for the data used varied from customer to customer, but in all cases covered a 12-month period. As stated in Section 1.1, Crystal Energy also provided half-hourly actual and predicted SMPs from 01/01/2008 - 01/06/2011. Table 2 shows the LWAP of the different customers, the total system LWAP for the corresponding time period, the ratio between the two and the time period over which these measures were calculated. In each case, the actual SMP was used for $Price_t$ in formulas (2) and (3).

Customer	Customer LWAP	Total system LWAP	Ratio	Time Period
Cold storage	46.74	54.36	0.86	1/12/2009 - 30/11/2010
Fast food	56.81	57.17	0.99	1/1/2010 - 31/12/2010
Hotel	56.14	57.17	0.98	1/1/2010 - 31/12/2010
Manufacturing	53.62	52.75	1.01	1/11/2009 - 30/10/2010

Table 2: Ratio between customer LWAP and total system LWAP using actual SMPs

From Table 2 we can see that the ratio between the cold storage facility's LWAP and the corresponding total system LWAP is 0.86. This is below 1, indicating, as we expected from Figure 16, that the cold storage facility is well suited to the tariffs offered by Crystal Energy. The corresponding ratio for the manufacturing company is 1.01. As this is above 1, it signals that this customer is not suited to flexible tariffs.

The ratio between the customer LWAP and the total system LWAP for the fast food restaurant and the hotel are 0.99 and 0.98 respectively. These show that both these customers are suitable to the tariffs offered by Crystal Energy. However, as both ratios are much closer to 1, the benefit would not be as great as that enjoyed by the cold storage facility.

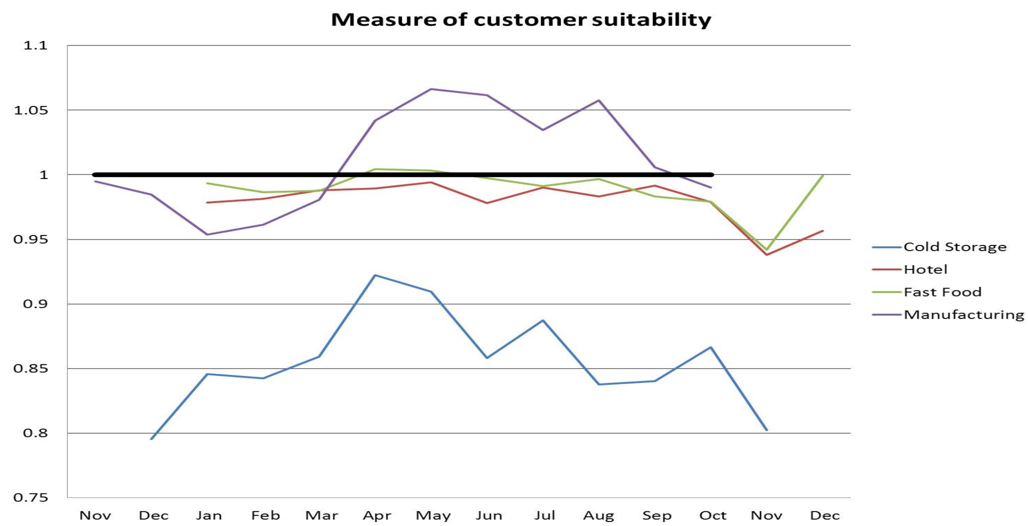

Figure 18: Ratio between customer LWAP and total system LWAP

In Section 2 we saw that the level of prices and the number of spikes varied across the different months of the year. As a result the group also decided to analyse LWAPs for each customer, on a month-by-month basis. Figure 18 shows the ratios of the customer LWAP to the LWAP of the total system, calculated for each month in the time period of the data provided. The ratio between the cold storage facility's LWAP and the total system's LWAP (blue line) remains consistently below 1 across the time period. This indicates that this customer is suitable for the tariffs offered by Crystal Energy at all times of the year and agrees with the ratio calculated for the overall time period.

Likewise, the ratio between the Hotel's LWAP and the total system's LWAP (red line) is consistently below 1, signalling that they are also suited to the tariffs offered by Crystal Energy throughout the year. However the Hotel's ratio remains much closer to 1 than that of the cold storage facility. This implies that the benefits of switching to Crystal Energy are consistently much greater for the cold storage facility.

Figure 18 also shows that the ratio between the fast food restaurant's LWAP and the total system's LWAP (green line) is below 1 for all except two out of the twelve months. This indicates that this customer is generally suited to the tariffs offered by Crystal Energy but there are times when they may not be. It also corresponds with the ratio calculated over the total time period for this customer (see Table 2). Note that the overall ratio does not capture the fact that this customer may not be suitable for the flexible tariffs at certain times of the year.

The ratio between the manufacturing company's LWAP and the total system's LWAP (purple line) varies significantly across the time period of the data provided. In January and February of 2010 the ratio is ≈ 0.96. However after this, the ratio increases to ≈ 1.06. By November 2010, the ratio decreases again to below 1, suggesting this cycle may be seasonal. These figures indicate that this customer is, at times, clearly suitable for the tariffs offered by Crystal Energy, while at other times they are not. This is in contrast to the somewhat ambiguous ratio calculated over the entire time period of the data provided (see Table 2) which is very close to 1.

4.2 Robustness of measure to changes in price

Customer	Ratio using predicted SMPs	Ratio using actual SMPs
Cold storage	0.87	0.86
Fast food	1.00	0.99
Hotel	0.99	0.98
Manufacturing	1.02	1.01

Table 3: Ratio between customer LWAP and total system LWAP using both predicted and actual SMPs

To test the robustness of this measure to changes in price the group undertook the same analysis but with predicted SMPs instead of actual SMPs. Table 3 shows the ratios between the LWAP of the different customers and the total system LWAP using both the predicted SMPs and the actual SMPs. As we've seen in Section 2 there can be large discrepancies between actual and predicted SMPs. However, despite this, Table 3 shows there are relatively small differences between the ratios when the two different types of prices are used.

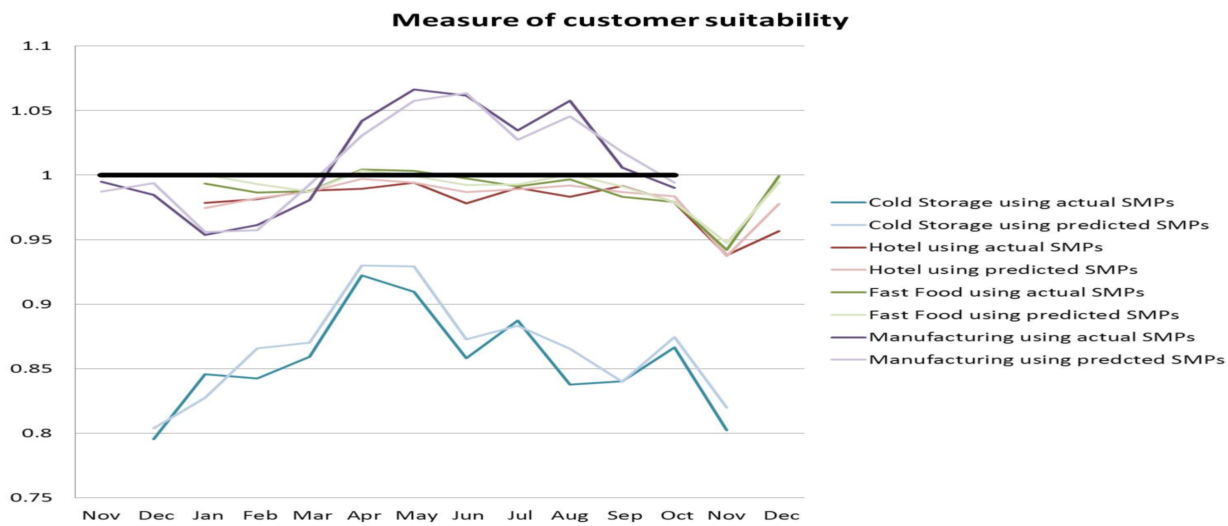

Figure 19: Ratio between customer LWAP and total system LWAP using both predicted and actual SMPs.

Figure 19 shows the ratios between the customers' LWAP and the LWAP of the total system for both actual and predicted SMPs across the different months of the data provided. The ratios calculated with the actual SMPs are in the darker colours while those calculated using the predicted SMPs are in the corresponding lighter colours. This graph shows that there are only small differences between the ratios calculated using the two different types of prices. This is consistent with the ratios calculated over the entire time period of the data provided, see Table 3. As a result, it can be said that this measure of customer suitability is robust to fluctuations in price.

5 Summary

Crystal Energy presented two problems to the Study Group: the question of price predictability, and the development of a metric for customer suitability. The first problem was solved by applying the statistical analysis of Section 2, which produced a measure of expected error for each half-hour interval in each month of the year. This information was added to the ex-ante price time series, giving a clear visual presentation of the risk associated with moving load to certain time intervals (see Figures 8 - 11). Further information is also provided by the probability intervals for spike occurrence, see Figures 13 - 15. The solution of the second problem is described in Section 4, with the metric giving very useful information on the suitability of various customers for Crystal Energys pricing tariff, as shown in Figures 18 and 19. More detailed information on the seasonal variation of the metric, as well as its robustness to price prediction errors, was also provided.

Acknowledgments

All contributors would like to thank Paddy Finn for introducing the problem and assisting in answering questions and formulating the solutions. We also acknowledge the support of the Mathematics Applications Consortium for Science and Industry (www.macsi.ul.ie), and the Sustainable Electrical Energy Systems research cluster, funded by Science Foundation Ireland through grants 06/MI/005 and 09/SRC/E1780, respectively.

References

[1] "Solver Choice in the SEM - A Comparative Study of Lagrangian Relaxation vs. Mixed Integer Programming ," SEM Market Studies Publications.

[2] "Market Incident Report September 20th 2010," SEM Market Studies Publications.

[3] http://www.macsi.ul.ie/2/esgi70/ESGI70_proceedings.pdf, page 97.

[4] Jabłońska M, Mayrhofer A, Gleeson JP, “Stochastic simulation of the uplift process for the Irish electricity market,” Mathematics in Industry Case Studies Journal, **2**, 88-110 (2010).