www.concept.co.nz

An approach to D+1 allocation

Prepared for Gas Industry Company

18 March 2015

About Concept

Concept Consulting Group Ltd (Concept) specialises in providing analysis and advice on energy-related issues. Since its formation in 1999, the firm's personnel have advised clients in New Zealand, Australia, the wider Asia-Pacific region and Europe. Clients have included energy users, regulators, energy suppliers, governments, and international agencies.

Concept has undertaken a wide range of assignments, providing advice on market design and development issues, forecasting services, technical evaluations, regulatory analysis, and expert evidence.

Further information about Concept can be found at www.concept.co.nz.

Disclaimer

While Concept has used its best professional judgement in compiling this report, Concept and its staff shall not, and do not, accept any liability for errors or omissions in this report or for any consequences of reliance on its content, conclusions or any material, correspondence of any form or discussions, arising out of or associated with its preparation.

No part of this report may be published without prior written approval of Concept.

Contents

1 Introduction

The potential introduction of a 'Market-Based Balancing' (MBB) mechanism on the Maui pipeline involving daily cash-out of any user imbalances over allowed tolerances, gives a greater requirement on shippers to manage their positions on a day-to-day basis.

However, at present, shippers have relatively poor visibility of how much their consumers have used (or have been deemed to have used through the allocation process) until quite some time after the gas trading day:

- For their non-TOU consumers (allocation groups 4 & 6), the Initial allocation doesn't occur until the 4th business day of the next month i.e. 'M+1'.
- For their TOU customers who don't have telemetry (allocation group 2) there can be a delay of many days or even weeks before meters are interrogated.

Accordingly, there is a need for a so-called 'D+1' allocation process which will occur the day after the trading day. Such a close-to-real-time allocation will give shippers information as to the extent to which their allocated gas for the previous day differed from their nomination, and thus the extent of their imbalance. This information will help shippers make improved nominations for the following day in two respects:

- Firstly, it will help them to correct for any imbalance from the previous day. For example, if they over-nominated by 2 TJ on the previous day, they can correct their running imbalance by reducing their nomination for the following by 2 TJ.
- Secondly, irrespective of being able to correct for previous days' imbalances, improved information on consumption for the previous day may help shippers make more accurate estimates for the following day.

An example of how a shipper might use the D+1 allocations to balance their ROI through a month is show in Appendix A.

Through helping individual shippers manage their positions, the D+1 allocation process should reduce the extent to which the Maui pipeline as a whole gets into a situation of imbalance which exceeds allowed tolerances.

This document describes a proposed process for producing a D+1 allocation prediction for retailers on the BoP, North and SKF pools on the Vector transmission network.

It is a fairly technical document, as it is assumed that the audience will be familiar with the current approach, and associated issues, regarding allocation and pipeline balancing.

2 Scale of the issue

Before we describe the D+1 algorithm we've developed, it is useful to remember what proportion of the daily gas volumes need to be allocated by D+1. The chart below shows gas volumes since October 2009 carried through the Vector transmission lines. Direct connect and bypass gas gates are shown separately; for the shared gas gates, volumes are broken down by allocation group.

For direct connect gas consumption and volumes associated with bypass gas gates, there is already information available on a daily basis, on the day after gas flow. Those volumes – about 65% of all gas delivered on the Vector pipelines – are not included in the D+1 calculations.

The focus of D+1 is on allocating gas volumes at the shared gas gates. Consumers in allocation group 1 have telemetry associated with time of use (TOU) meters, so their D+1 allocation should be straightforward. Group 1 ICPs account for about 15% of load on the Vector system.

That leaves groups 2, 4, and 6 for D+1 to allocate – a bit under 25% of volumes delivered on the Vector system.

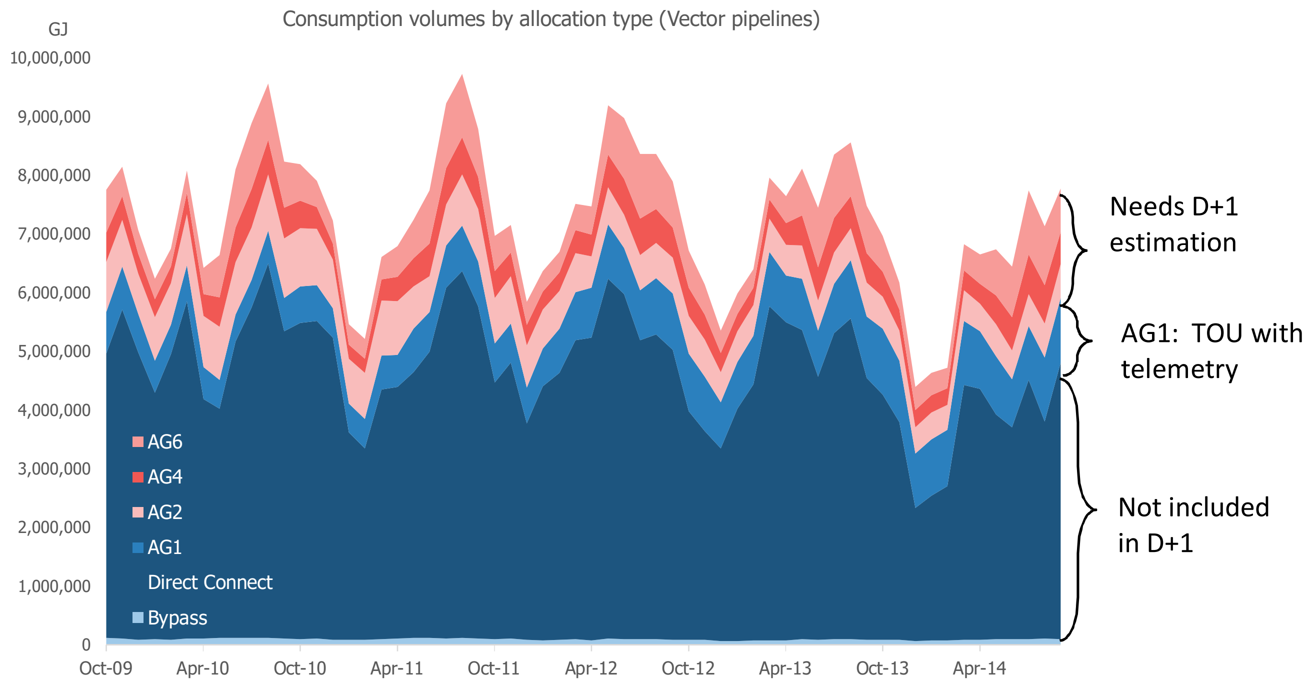

3 Description of D+1 allocation approach

3.1 High-level overview

Problem definition

The challenge in attempting to allocate gas to retailers soon after the trading day is that the majority of consumers' meters are not read until many weeks (and in some cases months) after the trading day.

The fact that Initial estimates may not be accurate is partially addressed for the purposes of purchasing wholesale gas through a series of wash-ups. Thus, after the ***Initial*** 'M+1' (one month after the consumption month) allocation, there are wash-ups following the ***Interim*** 'M+4' and the ***Final*** 'M+13' allocations.

While this wash-up process ensures that downstream volumes are eventually attributed to the correct retailer (to the extent possible), it does nothing to assist in making accurate nominations during the consumption month in order to avoid exposure to balancing charges.

Accordingly, for D+1 allocation it is necessary to develop techniques which can estimate as accurately as possible the consumption from retailers' customers using what limited information is available.

In this respect, the relevant information which should be known about a trading day in the hours immediately following it includes:

- Metered offtake at each Maui TP welded point.
- Metered offtake at each Vector gas gate.
- Allocation group 1 consumption at each Vector gas gate.
- The number of ICPs (customers) for each retailer at each Vector gas gate.
- The AUFG factor for non-G1M gas gates.

Additionally, there is historical information from the various allocation phases. Thus Initial allocation data will be known up to about 1 or 2 months into the past. Interim data will be known up to about 4 to 5 months into the past, and Final data will be known up to about 13 to 14 months in the past.

At a high level, the D+1 estimation approach uses all this information in a process which consists of the following steps:

1) Start with the metered delivered volume at each Vector gas gate
2) For the G1M gates estimate UFG for the gate using a regression model[1]
3) Subtract UFG-scaled group 1 consumption
4) Estimate the UFG-scaled group 2 consumption using another regression model
5) Subtract this estimated group 2 consumption to give a 'residual' volume at each gas gate representing the estimated total non-TOU consumption
6) Aggregate all residual volumes for all gates within a transmission pool (i.e. North, BoP, or SKF)
7) Allocating this pool's residual gas volume amongst retailers using a third regression model

This is illustrated diagrammatically in Figure 1 below.

[1] For non-G1M gates the AUFG factor is known with certainty.

Figure 1 - How gas volumes are split

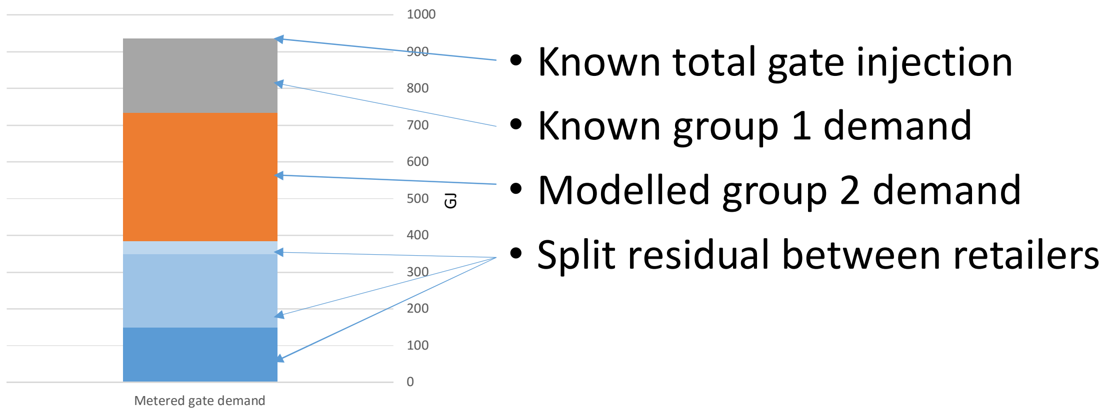

Regression Modelling

As indicated above, three regression models are used in the process. They are produced using "R", an open source statistical modelling and analysis package developed at the University of Auckland and used worldwide.

These models determine the relationships between independent parameters (other data that can be used to predict the value) and the dependent parameter (the value that is being predicted) based on historical data. Those relationships can then be used to predict the dependent parameter in the future, when only the independent parameters are known.

The best fit line on Figure 2 is an example of a very simple regression model with one independent variable.

Figure 2 - Simple regression model

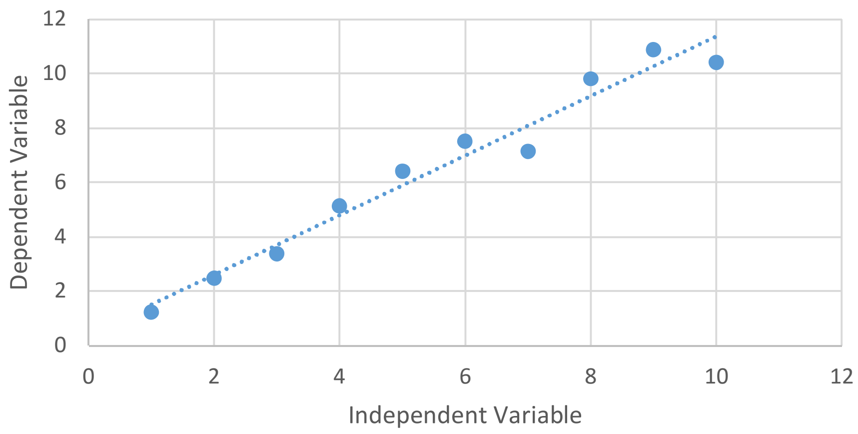

The regression models used in this D+1 estimation are *multi*variable regression models. Multiple independent variables are used to predict a single dependent variable. These models are described in more detail in the remainder of this section.

Different versions of the regression models are produced for each month, using the most up to date data that would have been available at the time.

3.2 Description of individual modelling steps

3.2.1 Aggregated offtake at each Vector gas gate

The amount of aggregate offtake at each Vector gas gate is known either in almost real-time (for gates with SCADA) or by the day after (for gates with telemetry metering). It is assumed that there are no issues with accessing this information when performing the D+1 allocation.

3.2.2 Calculating UFG

In the monthly allocation process, consumption at a gas gate is corrected to account for any difference between the total gas measured at the gate, and the aggregate of all gas metered at customers' premises. This difference is referred to as Unaccounted For Gas (UFG). To account for UFG, allocations to retailers are scaled up or down so that total allocated gas use is equal to the total gas measured at the gate.

There are two ways that UFG can be accounted for in the monthly allocation process. For Global 1-Month (G1M) gates, all allocation groups are treated equally, and submitted consumption volumes are scaled so that they are equal to the gas gate volume for that month.[2] Non-G1M gates are more complicated. For these gates, users in allocation groups 1 and 2 have their submissions scaled by the average UFG for the previous year. The remaining gas gate volumes are allocated to allocation groups 4 and 6 in proportion to their submissions.

For the purposes of producing D+1 predictions, the UFG factor for groups 1 and 2 at non-G1M gates will be known. However, for G1M gates, the UFG factors need to be estimated.

A regression model was developed to do this estimation, with a different version of the model for each gate. Each version has the same structure, but the relationships between parameters are different for each gate. The model uses three parameters to produce the estimate:

- UFG at the gate from the previous month.
    - UFG is often correlated with the previous month's UFG.
- The month of year.
    - UFG is often seasonal.
- A time factor.
    - UFG values can 'drift' over time.

This is quite a simple model – the only data required to produce it is the UFG from the previous month. However, it performs much better than simply assuming a UFG of 1 for each gate.

3.2.3 Group 1 consumption

The amount of gas used by each group 1 ICP should be reported soon after the end of the trading day. For non-G1M gates, the allocation of group 1 ICPs is straightforward, as it is just the daily consumption multiplied by the AUFG factor. However, for G1M gates an estimated UFG has to be used, which means that the group 1 D+1 allocation will change with successive allocation runs, as more accurate information on mass market consumption volumes will be available.

Gas meters measure the volume of gas, which then needs to be corrected to standard conditions and multiplied by the calorific content to produce an energy value. In order to get the data

[2] Although this is the current method for dealing with G1M gates, this is not how they were allocated historically. For the D+1 estimation exercise, historical allocations have been recalculated using the method used for YTD September 2014. This means that historical allocations may differ from those used in the model for some gates.

necessary to supply D+1 allocation information in a timely manner, it may be necessary to use fixed or estimated correction factors to produce an estimated energy value. This estimate would then be corrected at a later allocation stage.

3.2.4 Group 2 consumption

Although group 2 ICPs are metered daily, their usage is not known until the end of the month because they don't have telemetry. Accordingly, it is necessary to estimate the consumption of each group 2 ICP. This is achieved for the D+1 process using another regression model.

Each of the approximately 250 ICPs is modelled with its own individual model. This is preferable to modelling group 2 consumption at a retailer level, because a group 2 ICP switching from one retailer to the other produces large changes in group 2 retailer consumption.

For most ICPs, the analysis has revealed that the best parameters for modelling the consumption for a group 2 ICP for any given day are:

- Previous month's consumption
    - ICP consumption is often highly correlated with previous month.
- Month of year
    - ICP consumption is often seasonal.
    - An alternative way to include seasonality is to use consumption from the corresponding month in the previous year. However, including "month of year" as a parameter for the model produces better results.
- Injection at the gate
    - ICPs can sometimes make up a large proportion of the total gas gate injection volume. If this is the case, then gate injection and ICP consumption will be highly correlated.
- Business day or non-business day
    - Many ICPs have lower consumption on non-business days.
    - There was no benefit found in further differentiating days into day of the week (Monday/Tuesday/etc).

The model uses a weighting so that more recent values have more impact on the modelled fit.

Newer ICPs (less than a year old) don't have sufficient consumption history to use 'month of year' as a prediction parameter. For these ICPs, the model doesn't include seasonality. For even newer ICPs, the previous month's consumption can be misleading, since it can change substantially during commissioning. For the newest ICPs, previous month's consumption is not used in the model.

Ideally, new group 2 ICPs would directly supply the allocation agent with their estimated expected consumption. This would be desirable for the first year's operation. Similarly, any group 2 ICP that unexpectedly ceases using gas should notify the allocation agent.

Some ICPs switch between group 1 and 2 over time. To account for this, a model is produced for the entire time period for any ICP that is ever in group 2. When the ICP is in group 2, results from the model are used. When the ICP is in group 1, it is assumed that its actual consumption is known for the D+1 allocation.

3.2.5 Allocating Residual Gas

Once consumption for groups 1 and 2 has been estimated, the difference between the sum of that consumption and the gas gate injection quantity is the daily residual gas volume.

In the current monthly allocation process, retailers are allocated a share of the daily residual gas based on their share of the total non-TOU submissions at that gate.

The proposed approach for the D+1 process uses historical retailer submissions as one of a number of pieces of information feeding into a third regression model.

This model is slightly different from the other models. The other models predict absolute values: either the UFG value for a gate, or consumption for a group 2 ICP. This model predicts the *share* of the residual gas volume that is attributable to each retailer for the entire pool. The predicted proportion is then applied to the sum of the residual gas volumes for all gates within a pool to produce allocations.

A different version of the model is produced for each pool. The parameters used for these models are:

- Month of year
  - Retailer share is seasonal.
- Previous proportional allocation for the pool for the most recent historical month, based on retailer submissions
  - Retailer share tends to not change substantially over a short period of time.
- *Proportion* of non-TOU ICPs serviced by each retailer
  - Some of the change in retailer share that does occur can be estimated by looking at changes in ICP numbers. The number of ICPs is available in near real-time.
  - The proportion of ICPs is better at estimating the proportion of residual volume than the number of ICPs.
- Retailer
  - Each retailer has its own mix of customers, which behave slightly differently

This model also uses a weighting so that more recent values have more impact on the modelled fit.

It is important to realise that because residual gas volumes are allocated based on monthly submissions in the monthly allocation process, the day of week or business day/non-business day does not have any effect on how the residual gas volumes are allocated. The model developed for D+1 is consistent with the current method of allocating gas.

The modelled proportions sum very closely to but not equal to 1, so they are scaled so that they do sum to 1.

As detailed in the second main bullet point above, one of the key inputs to the model is the submissions from retailers. These are used to determine the most recent 'known' retailer share, which is then tweaked as appropriate. Three different types of historical data were investigated to give an indication of which option may be best suited for the D+1 process.

1. A "one month behind" option, which uses the previous month's Initial allocation. This is the most up to date data that would be available for most days of the month. However, because the Initial submission for a month doesn't happen until 4 business days into the next month, the D+1 process would not have access to "one month behind" submissions for the very start of a month
2. A "two months behind" option, which uses the Initial allocation from two months previously. This is less up to date, but strictly speaking it is the most recent one that would be available with current processes given the four day delay in getting the Initial allocation.
3. A "two months behind – Interim" option, which uses the *Interim* allocation from two months previously. This is a more accurate submission than would be available in practice, but it

gives an indication of what might be available if the retailers submitted their *Initial* allocation data later in the month, after the SADSV values and a greater proportion of meter reads are available. The actual performance of using a delayed Initial data set, would likely be between the accuracy of this “two months behind – Interim” and the “one month behind” option.

4 Results

4.1 Allocations by pipeline pool

The D+1 allocation process described in this paper produces a non-TOU prediction for each pool, rather than by allocation gate. This “top down” approach allows the non-TOU consumption to be modelled better, since any vagaries at an individual gate will only make up a small part of the total non-TOU consumption in the pool. It is also consistent with the primary purpose of the D+1 exercise, which is to provide information to shippers that will enable them to make more accurate gas nominations and to manage mismatch more effectively.

Group 1 and 2 ICPs are modelled individually so a daily prediction could be produced for each ICP if it were desired.

4.2 As there can be multiple shippers owned by the same parent company, predictions are produced for the parent company (e.g. GEND, GENG and GEOL are included in the Genesis prediction and EDNZ and TRUS are combined).Overall results

The error of a D+1 allocation prediction is calculated as the difference between the prediction and the best known allocation value, taken from the Final allocation where available, or, for more recent months, the Interim. Results shown in this paper are presented for the 12 months to September 2014 (September 2014 being the final month for which Interim results were available at the time this analysis was undertaken.) Errors are presented as either a percentage or an absolute value. The percentage is the average absolute error as a proportion of the average best known allocation. How the absolute error is calculated is explained further in section 4.3.1 below.

Results for the three different options for retailer submissions are presented in most graphs, along with the ‘current’ error. This current error is the difference between the current Initial allocation and the best known allocation.

Showing this current error is useful to see how the D+1 estimation compares with the accuracy of the current Initial allocation process.

4.3 Results by group

Figure 3 – Average daily percentage errors, by group

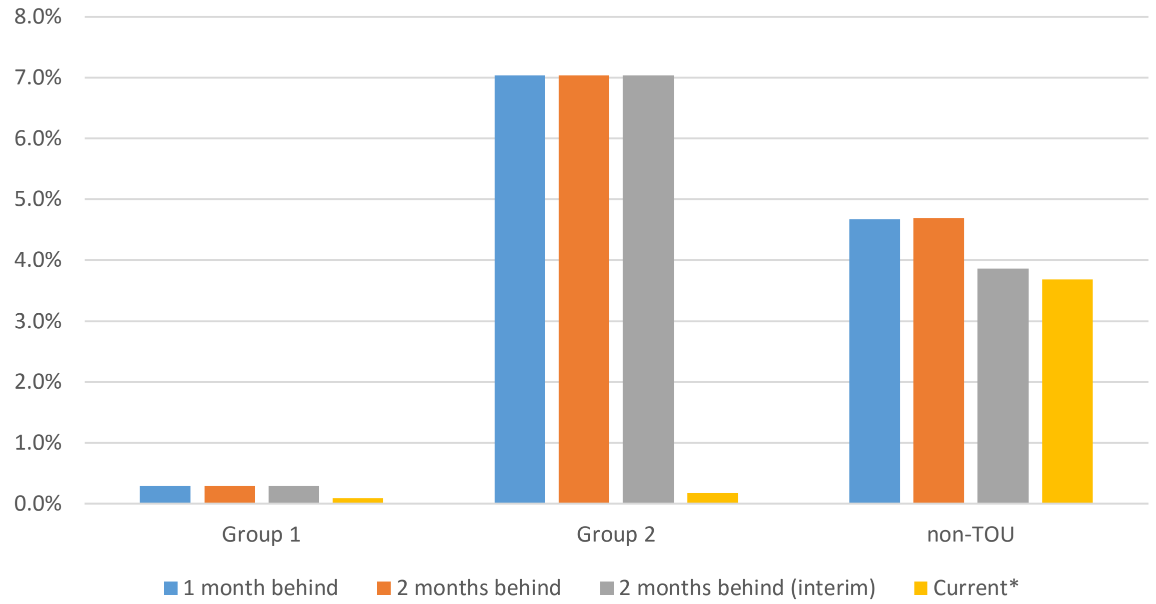

Figure 4 - Average daily GJ errors, by group

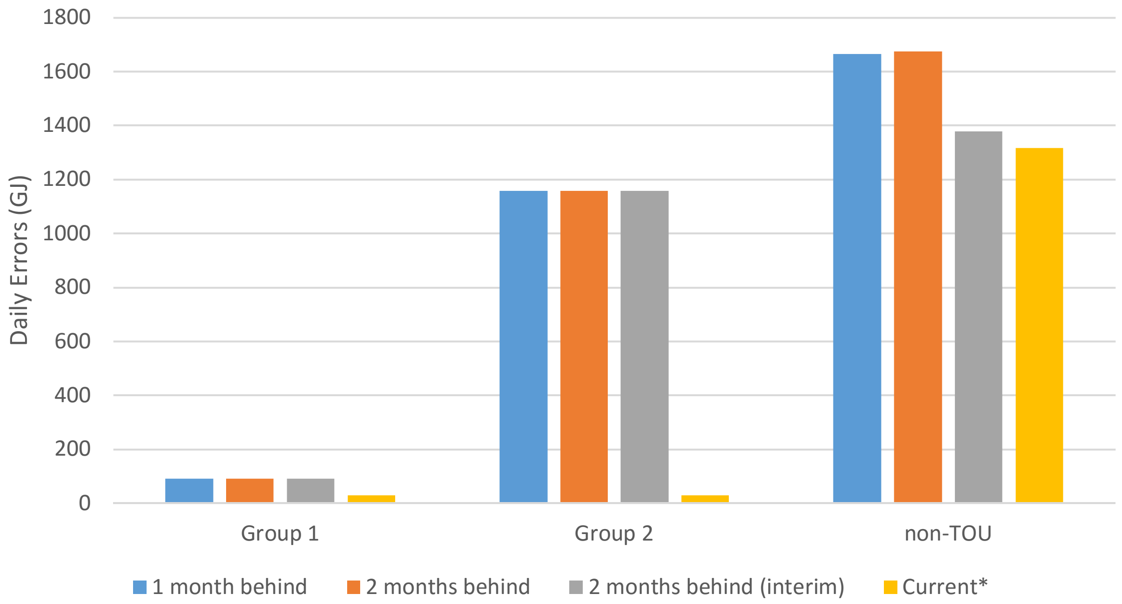

The errors for group 1 using the D+1 allocation process are small. The only source of error is in the UFG prediction for G1M gates. This can also occur in the current process (as shown by the small but noticeable error for group 1 "Current") when UFG values for a gate are revised in later allocations. There is an increase in group 1 errors for the D+1 process, but they are still less than 0.5%. For many pool and retailer combinations there is no error, as they do not have any group 1 ICPs at G1M gates.

Group 2 has the highest percentage error of the three groups. However, it is also the group with the lowest total consumption, so it has a smaller GJ error than non-TOU consumption. There is no difference between the 3 submission options for group 2, since submitted consumption does not typically change between Initial and Final submissions. Group 2 predictions are explored in more detail in section 4.4 below.

Non-TOU consumption is modelled well by the D+1 process, with the accuracy being only slightly worse on average (about 350GJ) than the current Initial allocation. This is, perhaps, unsurprising since retailer submissions are used to make the D+1 so their errors are related. The "2 months behind – Interim" option's error is even closer to the current error (60GJ worse, on average), which reflects the better accuracy of the Interim submissions. Non-TOU predictions are explored in more detail in section 4.5 below.

4.3.1 Further explanation of error calculations

For every day, a D+1 prediction is produced for each of the ~60 combinations of retailer, pool and allocation group (representing 3 pools, 3 groups and 5-7 retailers). The absolute difference between this and the best known consumption is the daily absolute error. To produce averages from this, each of the daily absolute errors are averaged over the entire time period to produce a daily average. There is one daily average for each of these combinations.

The average results shown in this section group these ~60 averages in different ways. Sometimes only a selection of them is shown (e.g. for non-TOU only averages, the other groups are ignored). In Figure 4 above, the errors for the ~60 combinations have been grouped by allocation group. So there are ~20 combinations that make up the values shown for each allocation group.

The average daily error for group 2 in Figure 4 above is about 1200GJ. This value is just the sum of the errors for the ~20 combinations. Figure 5 shows how this sum of errors is combined to make the total group 2 error. Some examples of retailer and pool combinations are labelled[3].

[3] Figure 5 is just intended to demonstrate how the average errors are calculated, and not to analyse the errors directly, which is done later.

Figure 5 - Breakdown of average group 2 error for a day

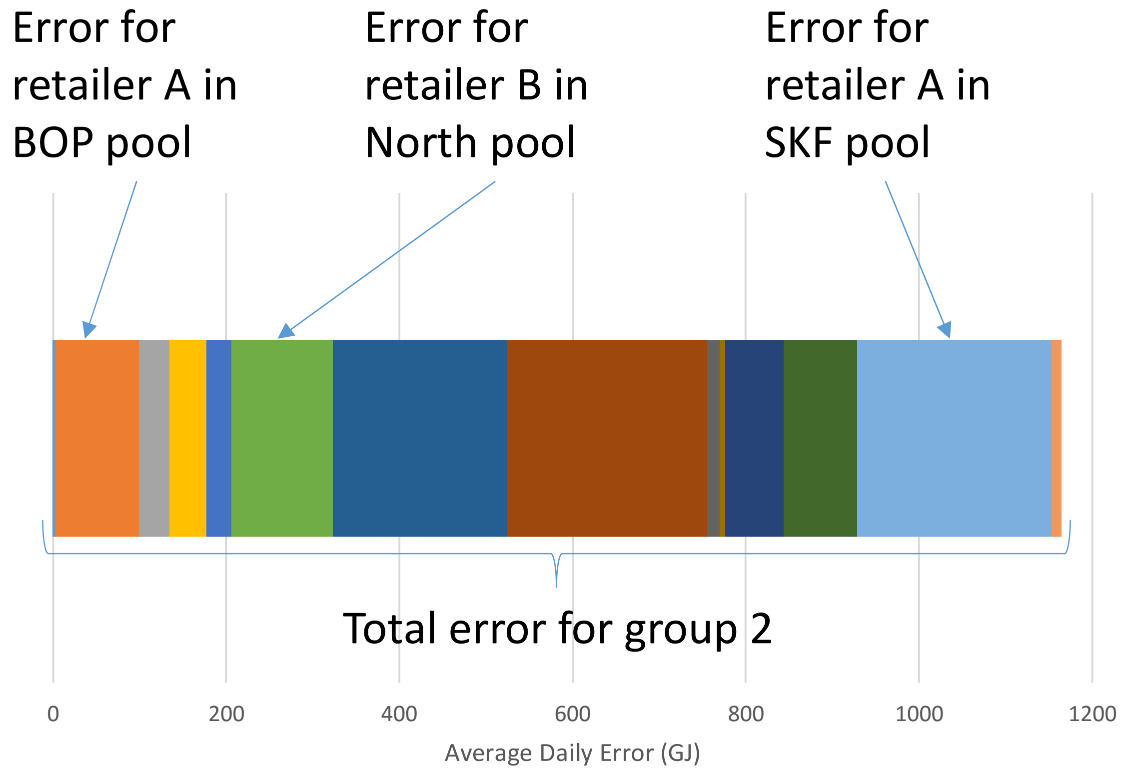

Total allocated volume is known exactly for a pool, because gate injections are metered, so the errors sum to zero at a pool level. However, it is the absolute errors that are used to calculate the accuracy of predictions, and these don't sum to zero. For example, consider three retailers and a single allocation group in a pool with errors of -50GJ, -50GJ and +100GJ. The errors sum to zero, as they must, but the average absolute error for that pool would be 200GJ since each error is added together.

Percentage errors are the absolute errors divided by the amount of consumption in that grouping. For the example above, the total group 2 average daily consumption is about 16,000GJ, so the percentage error is about 1,200/16,000 or 7%.

4.4 Group 2 examples

Group 2 consumption is often highly variable and less predictable than mass market non-TOU consumption, yet it is not reported daily like group 1. Some group 2 ICPs can be predicted well, either because they make up a large proportion of the total gas gate deliveries (and thus the telemetry metered gate volumes are a good indicator of the ICP's consumption), or because their consumption varies predictably by time of year or day of week. However, many group 2 ICPs cannot be predicted well because their consumption varies for other reasons. Unfortunately, absent moving these group 2 ICPs to telemetry reporting, it is difficult to see how the prediction could be substantially improved given the information available at D+1[4].

Figure 6 shows both the absolute and percentage error for all group 2 ICPs. Clearly some are much worse than others. The 50 largest (out of 238) ICPs are labelled 'Large' with the others labelled 'Small'.

[4] The effect of moving some of these ICPs to group 1 is investigated below.

Figure 6 - Absolute and percentage error for group 2 ICPs

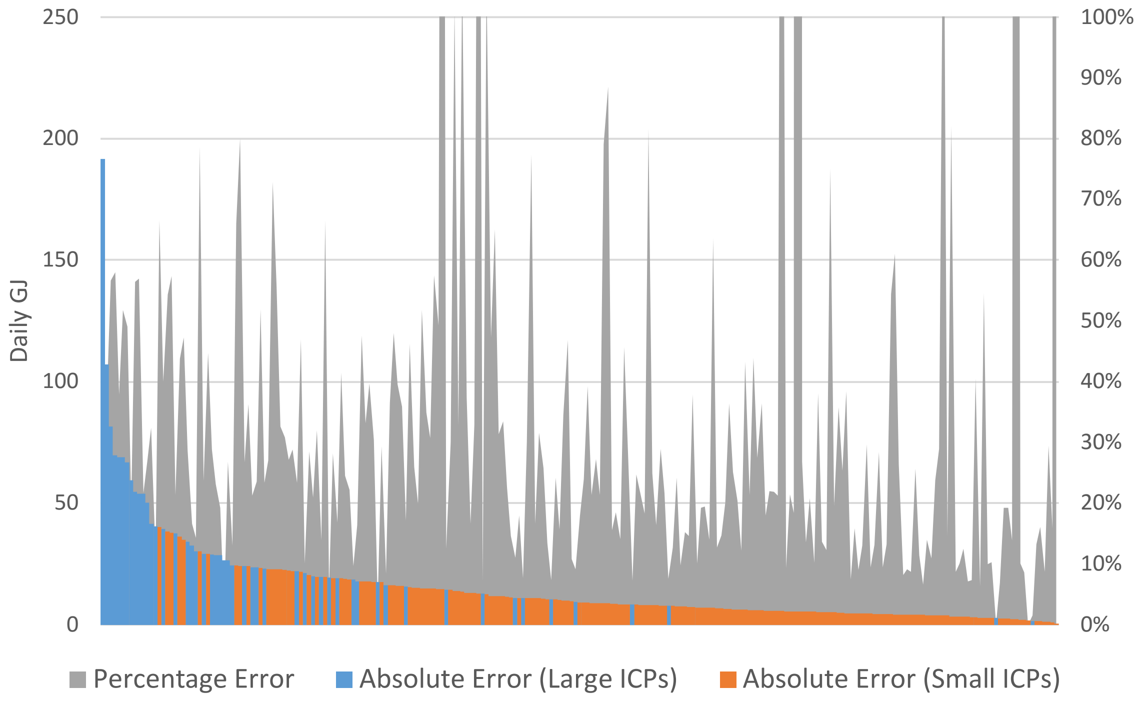

The worst (leftmost) ICP in the figure above is an ICP that started midway through 2014. It is very hard to predict consumption for this particular ICP, since there is limited historical data. This is a good example of a new ICP for which advising the allocation agent of average expected use would yield large improvements in accuracy.

Many of the other high error ICPs are simply highly variable, and reasonably large in size – although small compared to the total volumes for the gas gate where they are located. These ICPs are the type that might benefit the most from using telemetry and being moved to group 1.

Although all the worst performing group 2 ICPs are "large", there are also large ICPs that are easily predicted. This suggests that a GJ threshold for moving to group 1, while effective, would target more ICPs than would strictly be necessary.

4.4.1 Examination of potential impact of switching ICPs to group 1

Switching the 50 largest group 2 ICPs to group 1 by installing telemetry would reduce the total group 2 error by about 50%. Additionally, since this would reduce the error in predicting residual gas volumes, non-TOU consumption would also be predicted more accurately.

The effect on system error of switching the worst ICPs is shown below. Only the "one month behind" option is shown, but the effect on other options is similar.

Figure 7 - Effect of switching high error group 2 ICPs to group 1 on D+1 error

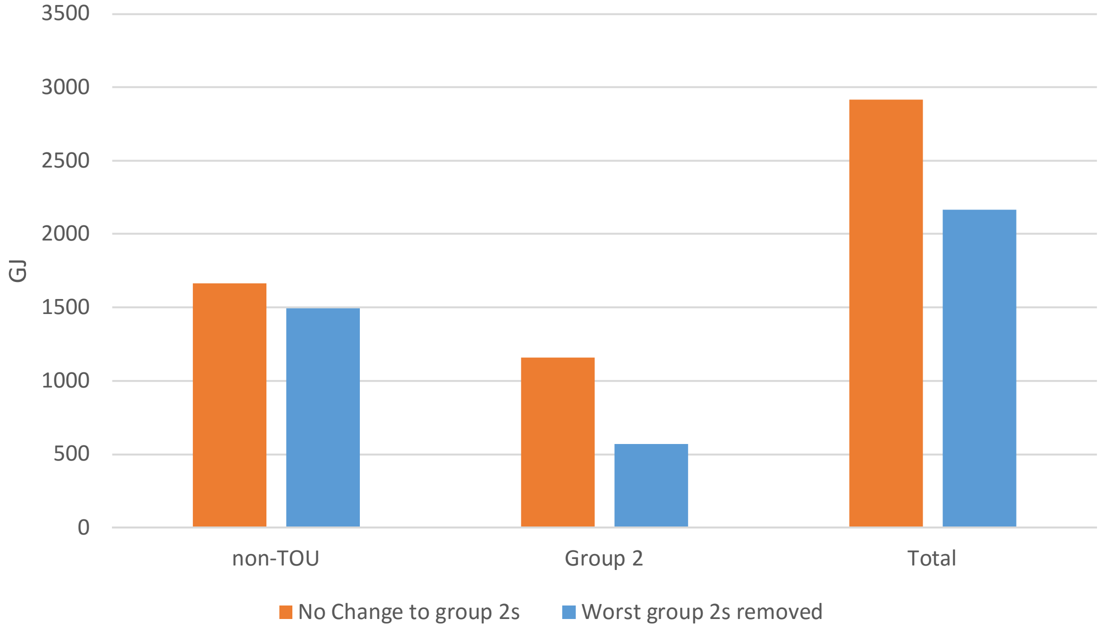

Obviously there are costs associated with installing telemetry and moving ICPs from group 2 to group 1. The net cost-benefit trade-off may or may not be beneficial, and has not been considered in this paper.

4.5 Results by retailer

In this section, and many of those that follow, the errors are not separated into the ~60 retailer, pool and allocation group combinations explained above. Instead, errors are only separated by retailer and pool. Errors are not distinguished by allocation group. This means that if a retailer is over-allocated for non-TOU consumption and under-allocated for group 2 consumption, then their total error for that pool is reduced. There are only ~20 retailer and pool combinations with allocation group removed.

In Figure 8 and Figure 9 the different pool errors are summed together, meaning that each retailer's error shown is the sum of their error from each pool.

Figure 8 - Absolute error by retailer for all consumption

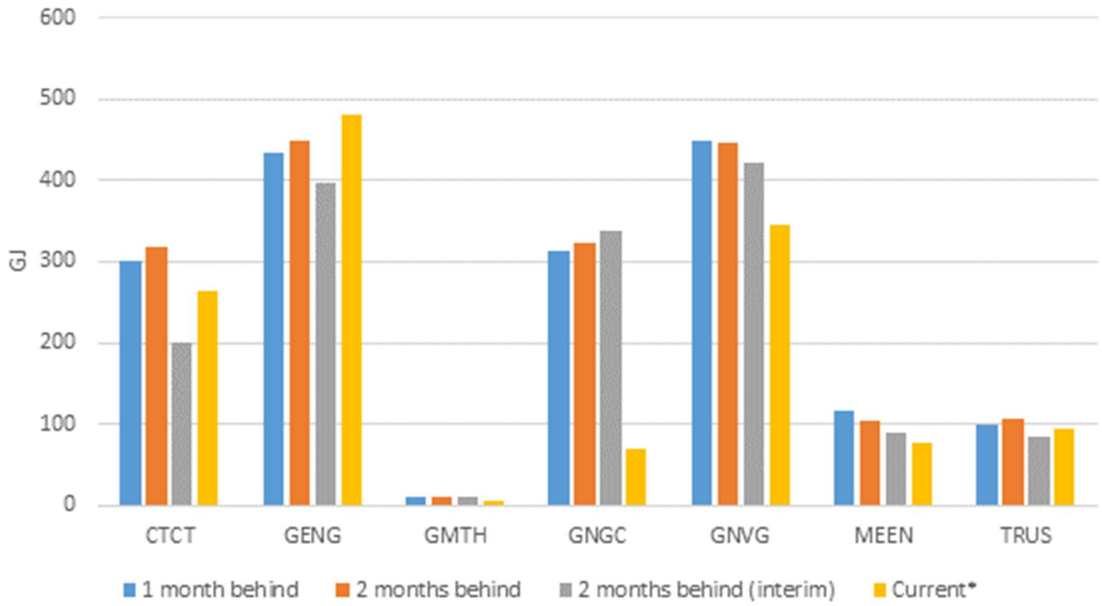

A retailer's contribution to total system error is highly influenced by the size of their allocated volumes. Larger retailers have larger errors in GJ terms, simply because they supply more gas overall[5].

There are also large differences between retailer's *percentage* errors.

Figure 9 – Percentage error by retailer for all consumption

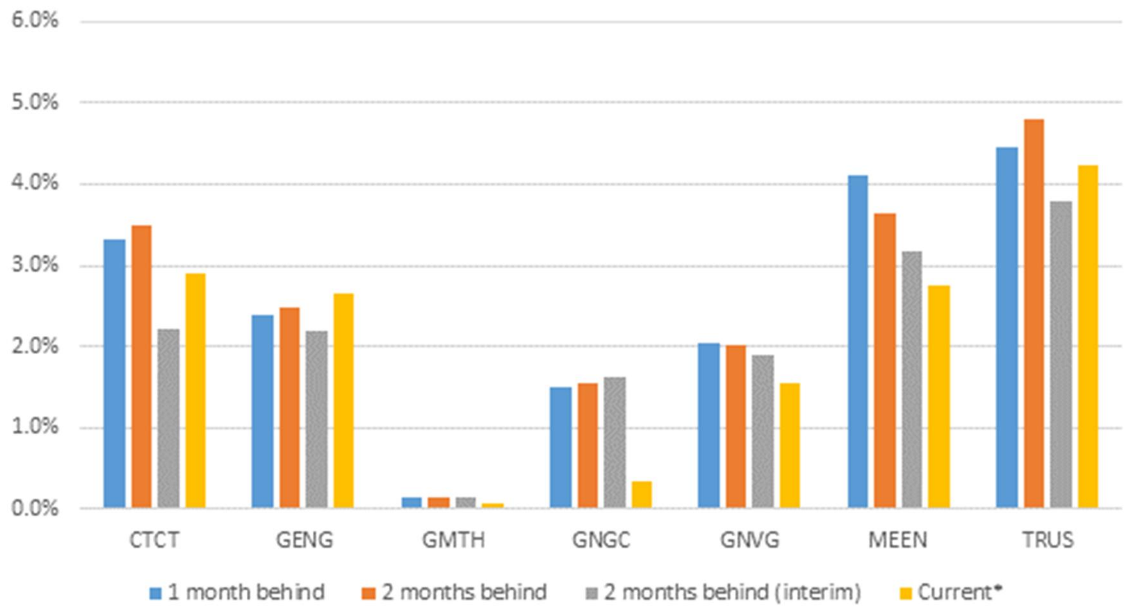

Group 1 consumption is much easier to predict than the other groups, so retailers with a large amount of group 1 consumption will do better. Figure 10 shows how each retailer's consumption is divided amongst the different allocation groups.

[5] GMTH is an exception to this, and is explained below

Figure 10 –How consumption is split between each allocation group for each retailer

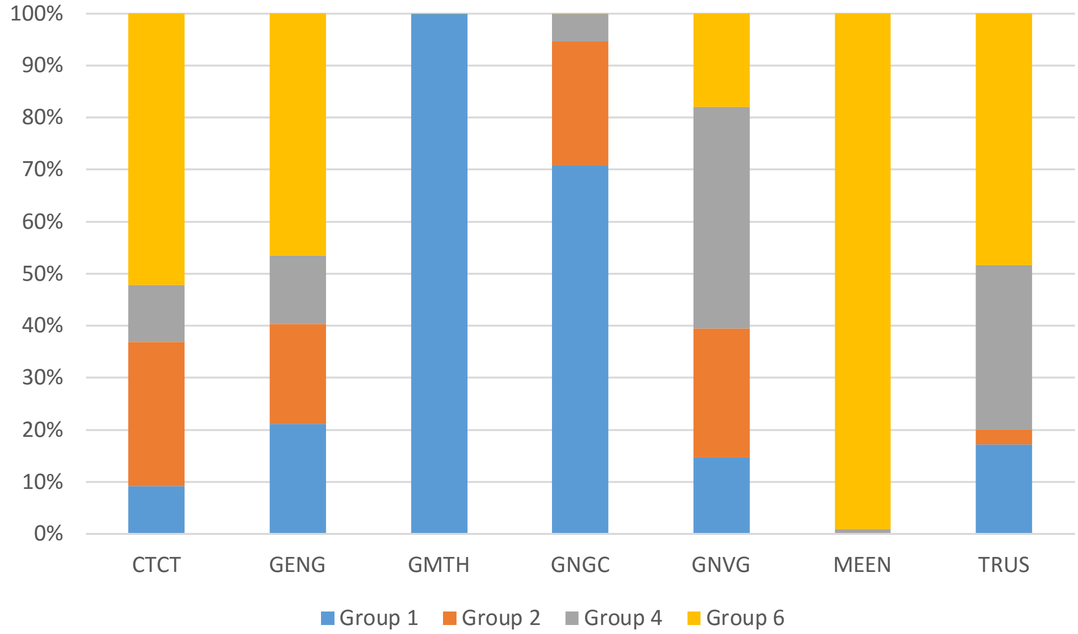

For most retailers, non-TOU makes up at least 50% their allocated volumes. However, GNGC and GMTH are different. Total non-TOU consumption makes up only 5% of GNGC's allocated volumes, and GMTH is almost entirely group 1. Accordingly, GMTH's percentage error is very small.

In light of this, GNGC is slightly surprising. Its percentage error is only slightly lower than most other retailers, despite 70% of its consumption being in group 1. This is because of how the modelling process works. The non-TOU allocations are calculated after taking away modelled group 2 consumption. So if the group 2 prediction is too high for a day, then the non-TOU prediction will probably be too low. In this way, when results are not separated by allocation group the predictions are somewhat self-correcting. If a retailer has both non-TOU and group 2 consumption then they will benefit from this. However, GNGC has a small amount of non-TOU consumption, so it does not. Similarly, MEEN and TRUS have a small amount of group 2 consumption and their errors are slightly higher than might be expected otherwise.

Average error is important, but a simple average can obscure some effects. A box and whisker plot is a more detailed way of presenting the error information[6].

In Figure 11, an error for each day and each pool is included in the box and whisker. Thus there are about 1100 values in each, made up of errors from 3 pools for one year. This is slightly different from how errors are presented in Figure 8, where they are summed between pools.

[6] Each box and whisker represents a retailer's errors for either the D+1 process (using the "one month behind" option) or the current allocation process. The whiskers range from the 5th percentile value to the 95th percentile value. The top box ranges from the 50th percentile to the 75th, whilst the bottom ranges from the 25th to the 50th

Figure 11 - Distribution of absolute errors for all pools for all consumption

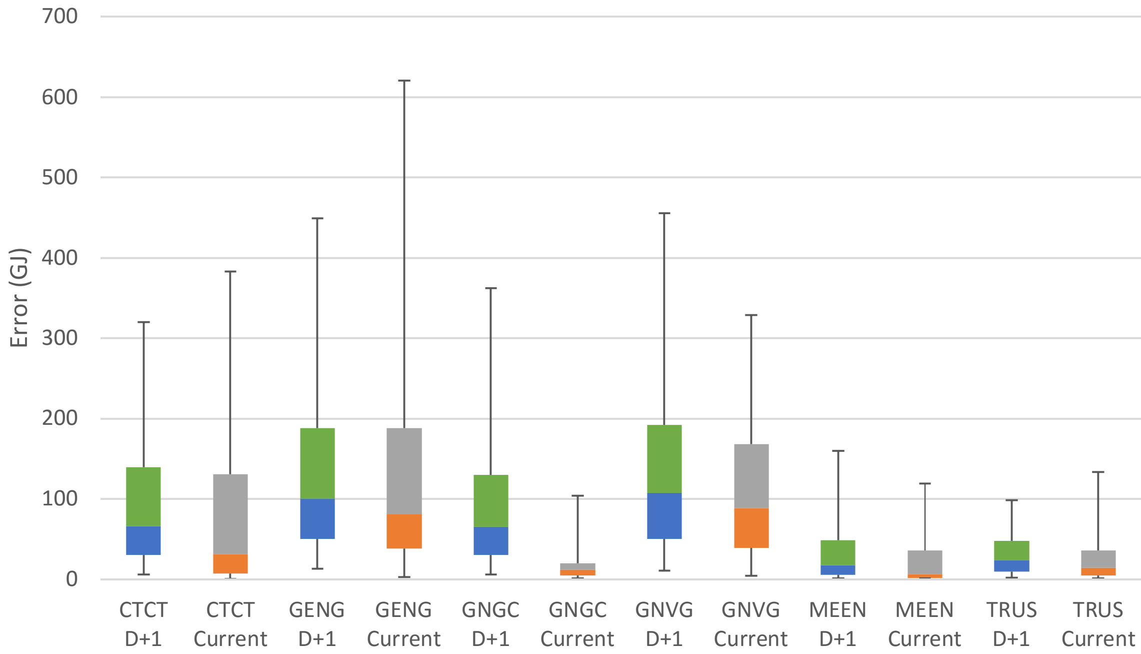

The top whisker for the GENG current allocation stands out. The mean and median errors for GENG are not significantly higher than other retailers, but large errors occur more often. This is replicated in the D+1 errors to a certain extent, but large errors do not happen quite as often. A similar reduction in large errors can be seen for CTCT, and TRUS.

This is encouraging for the D+1 process, as cashouts occur when the pool error is in excess of a tolerance. Reducing large errors may be more important than getting it right "on average".

GNVG and MEEN are slightly different. The average errors and large errors are higher for the D+1 allocation than for the current allocation. This may be because GNVG and MEEN currently read monthly, so their current allocation is more accurate.

GNGC's predictions are significantly worse. This is also the case when looking at average errors, and can be explained by similar reasoning[7].

4.5.1 Non-TOU results by retailer

It's also interesting to look at only non-TOU consumption, as the effect of group 1 and group 2 consumption is removed, so any differences are due to differences in the mass market models only.

For the non-TOU consumption graphs, only the non-TOU error is included in average calculations. That is, of the ~60 total combinations of errors available, only the ~20 that correspond to non-TOU error are grouped together for these graphs.

As for "all consumption" above, the absolute size of retailer's non-TOU error is mostly determined by the retailer's consumption volume. Larger retailers have larger error.Figure 12 and Figure 13 errors are calculated in the same way as for Figure 8 and Figure 9. Errors between different pools are summed to produce total retailer error.

[7] See above.

Figure 12 - Absolute error by retailer for non-TOU consumption

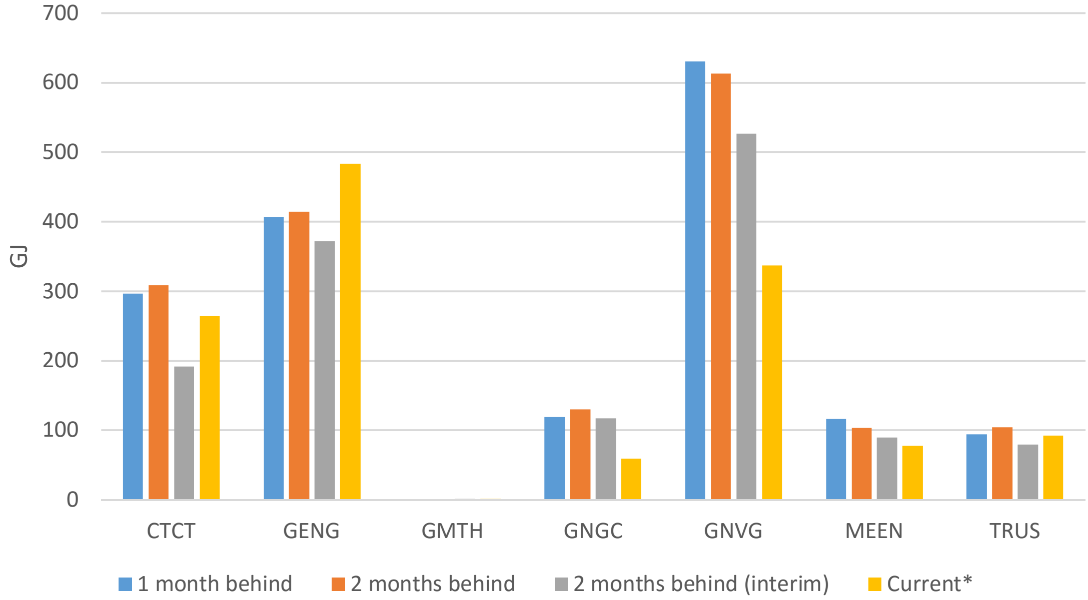

There are also differences in percentage error between the different retailers.

Figure 13 – Percentage error by retailer for non-TOU consumption

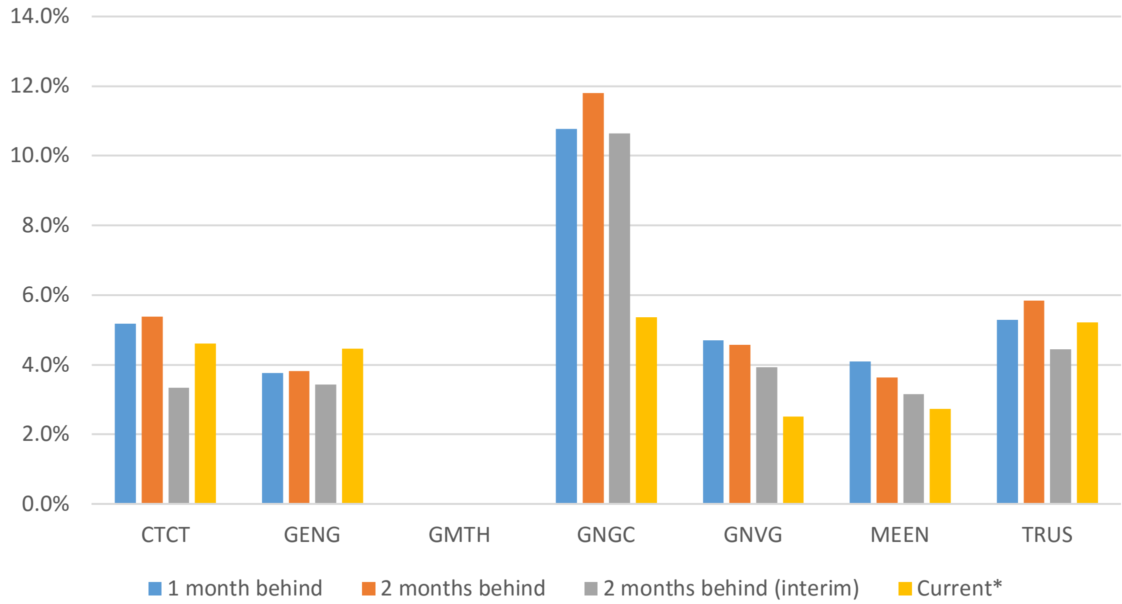

[8]

The improved accuracy of retailers who perform monthly meter reading is evident. GNVG and MEEN have noticeably lower current error than others. This improved accuracy in submissions probably also helps with producing more accurate D+1 predictions.

The D+1 model seems to produce smaller errors for GENG than the current initial allocation. This may be because GENG has a large number of group 6 ICPs whose consumption is relatively easy to predict.

[8] GMTH is not shown, as it has very low non-TOU consumption.

There is notably larger *percentage* error for GNGC, both in the current initial allocation and in the D+1 model results, compared to other shippers. This appears to be due to two main factors.

- There is slightly higher inaccuracy in Initial submissions (as indicated by their current error (the yellow bars) being higher than for other retailers).
- Most importantly, GNGC non-TOU consumption is made up of a small number of larger group 4 users[9]. The relative heterogeneity of group 4 consumption volumes is inherently harder to model than residential group 6 users who are more homogenous and predictable in their consumption patterns. Even though a different model is produced for GNGC (it doesn't assume that GNGC's ICPs are all group 6 users), a small number of bigger commercial users is harder to model than a large number of smaller residential users who all behave more similarly.

However, although GNGC has a large percentage error, its error is not large in absolute terms, since it has a small amount of non-TOU consumption.

The same effect can be seen to a lesser extent in GNVG non-TOU percentage error. GNVG has a large proportion of group 4 consumption, but it still has a significant amount of the more predictable group 6 demand.

4.5.2 Group 2 results by retailer

There are differences in group 2 errors between the retailers. However, it is hard to determine if any of the differences are indicative of real differences between retailers, or simply arise because the retailers supply gas to ICPs that are harder or easier to predict.

As such, and to avoid segmenting the results excessively, only system wide results are investigated for group 2. These are shown in Figure 14.

Figure 14- Absolute and percentage error for group 2 consumption for all pools

[9] See Figure 10.

The 3 major suppliers of group 2 consumption have similar percentage errors. They vary in absolute error roughly in proportion to their allocated volume.

The two other retailers stand out compared to the major 3 retailers:

- CTCT has significantly lower group 2 percentage error and slightly lower group 2 consumption. The combination of both these effects means that CTCT has significantly lower absolute error.
    - CTCT has a low number of group 2 ICPs (although some are large in size), and some of them make up a large proportion of their gate's injection. This means that the ICP's consumption can be forecast very accurately.
      Of course, this might not be a coincidence, as perhaps the worst performing group 2 ICPs have already been switched to group 1.
- TRUS has high percentage error, but low absolute error since they have very low group 2 consumption.

4.6 Results by pool

The size of the absolute errors varies greatly between different pools. As for the retailer calculations, this section does not separate error into allocation groups.

Figure 15 – Daily absolute error by pool for all consumption

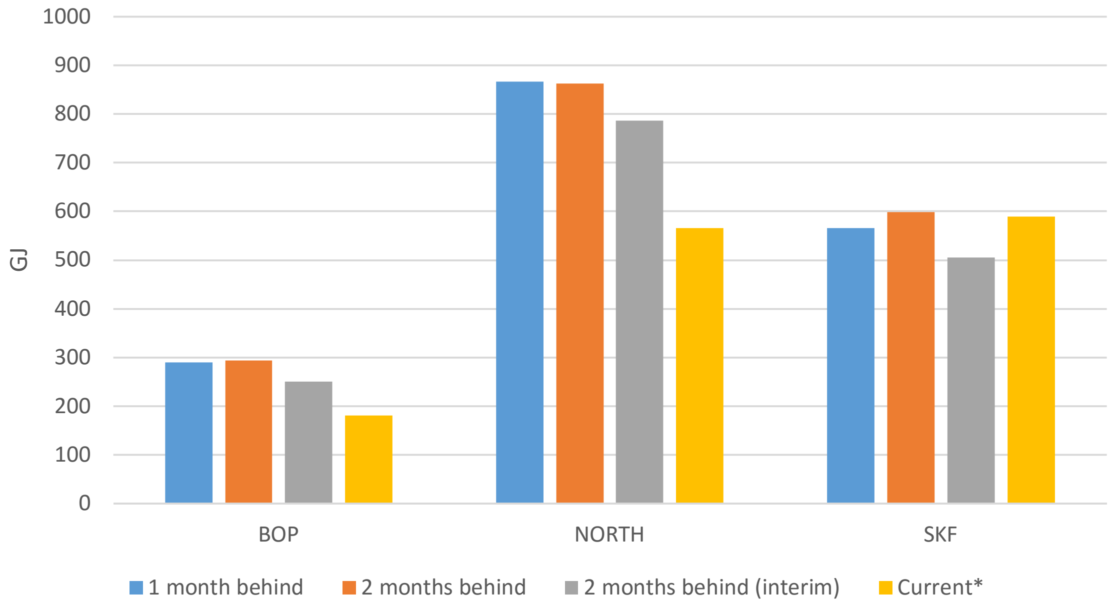

This is mostly because of the different amount of gas injected into each pool so that larger pools have larger errors.

Figure 16 shows the average daily *percentage* errors. There are only small differences in accuracy between the pools. The North pool has slightly more accurate allocations, both currently and when using the D+1 process.

Figure 16 - Daily average percentage error by pool for all consumption

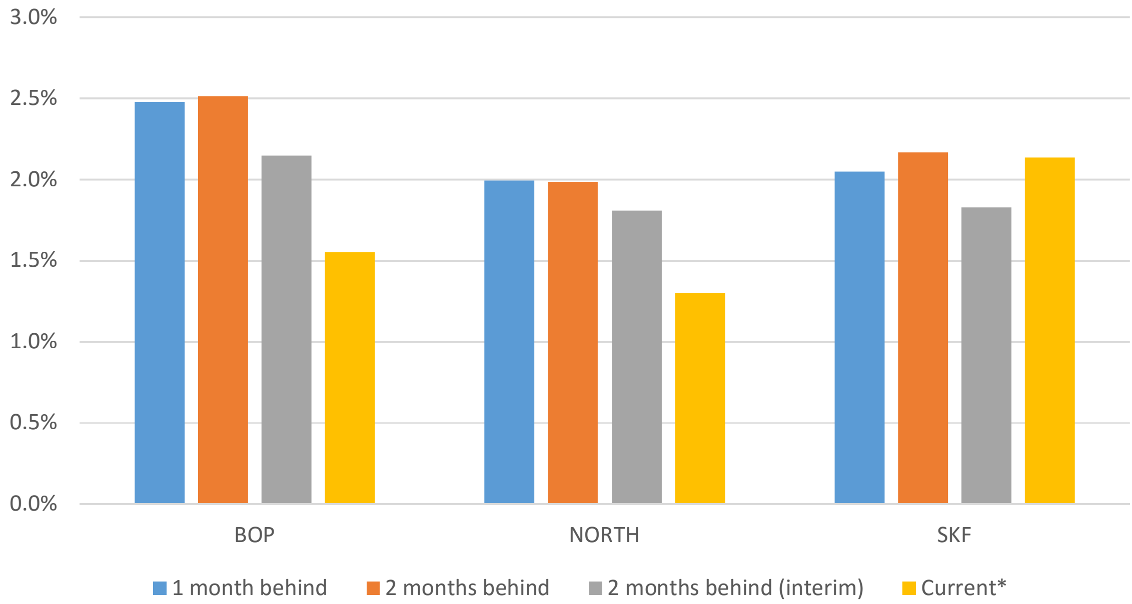

Similarly, for non-TOU it is the amount of consumption in each pool that is the main driver of differences. The BoP pool has the least error because it has the least non-TOU consumption.

Figure 17 - Daily absolute error by pool for non-TOU consumption

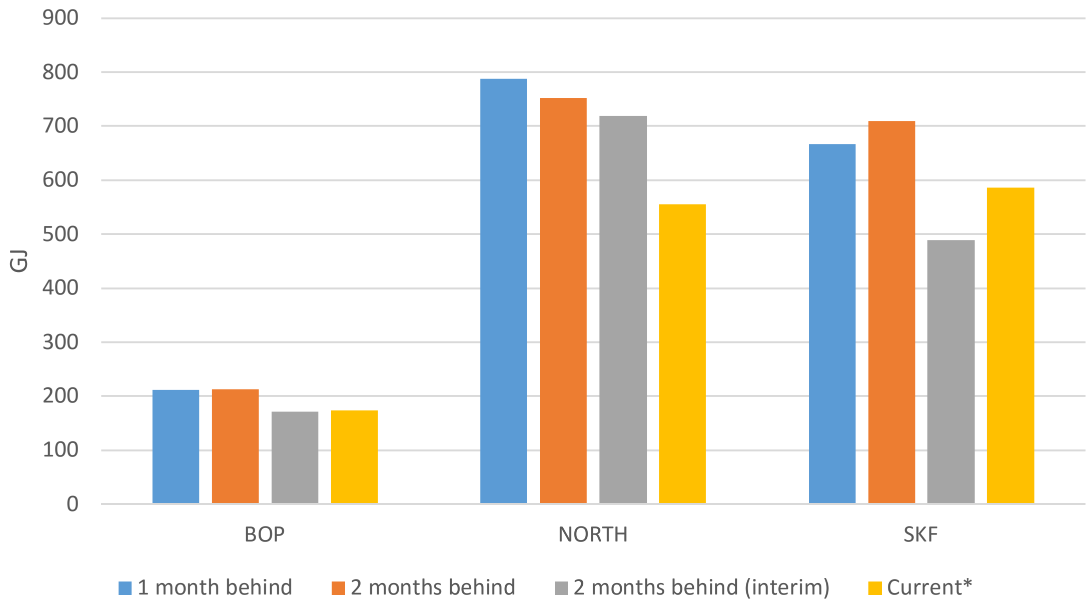

However, BoP has the largest *percentage* error, while North has the lowest. This difference occurs for both the current allocation and the D+1 allocation prediction.

Figure 18 - Daily percentage error by pool for non-TOU consumption

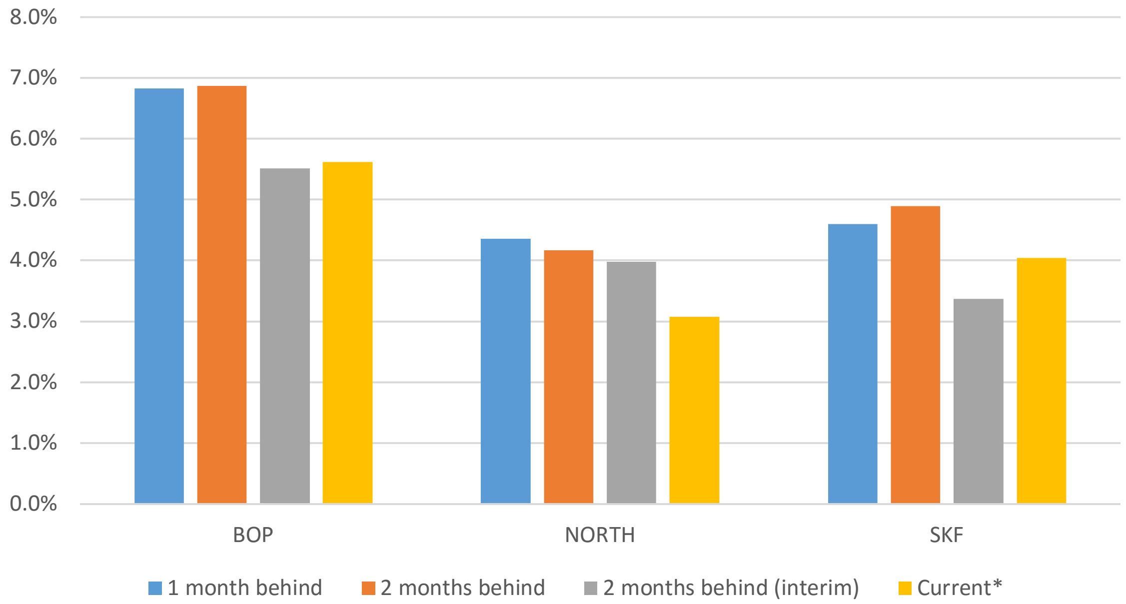

This may be because of the large shift of customers between GENG and TRUS in 2014 as shown in Figure 19.

Figure 19 - Share of non-TOU ICPs in BoP

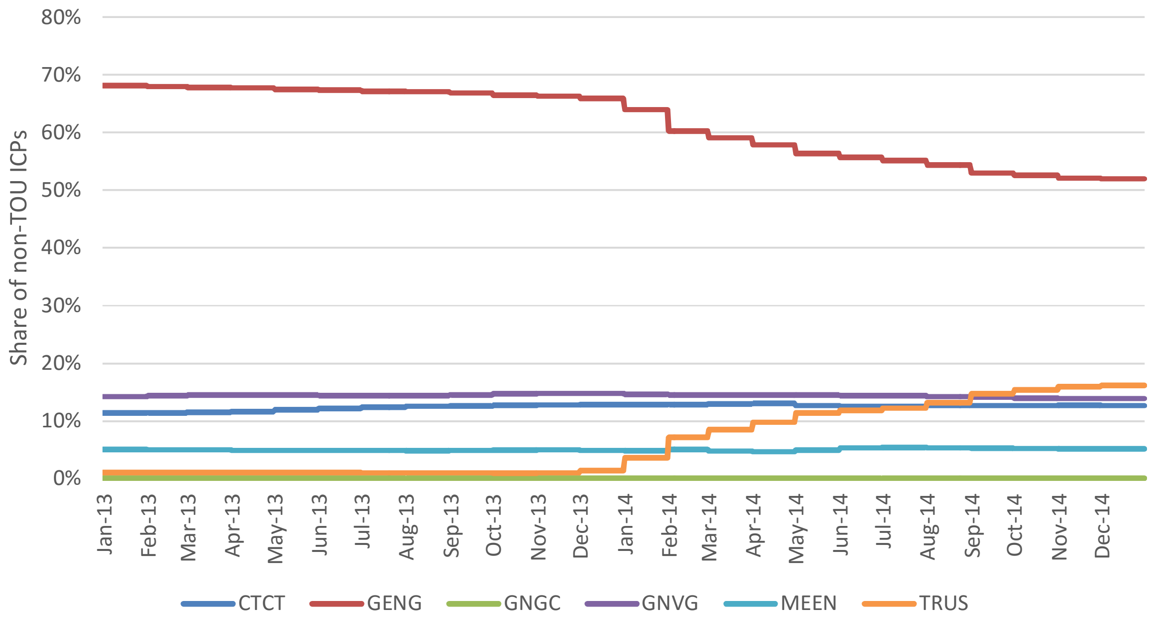

The model uses ICP numbers to adjust for any changes in customer base. However, this doesn't perfectly model the large changes in consumption experienced by TRUS. This may be because new customers behave differently from existing customers, so it is not only the number of customers that changes, but also the average customer consumption behaviour. Also, ICP numbers were only updated monthly for this exercise, and TRUS ICP share sometimes changed significantly in a month[10].

[10] ICP share more than doubled from Dec 2013 to Jan 2014

This same effect to a lesser extent can be seen in the North pool, where TRUS has also increased its market share.

TRUS has high percentage error in these pools, but since it makes up less than 7% of non-TOU consumption, this has little effect on total pool error.

Comparing Figure 15 and Figure 17 there appears to be very little difference between "all groups" error and "non-TOU only" error, which may be surprising. This happens because for the "all groups" analysis, the errors in different groups are combined, and as explained above, they tend to offset each other.

4.7 Further illustrations of retailer errors for each pool are shown in Appendix C. Results by day

Figure 20 - Daily error for each retailer for all allocation groups

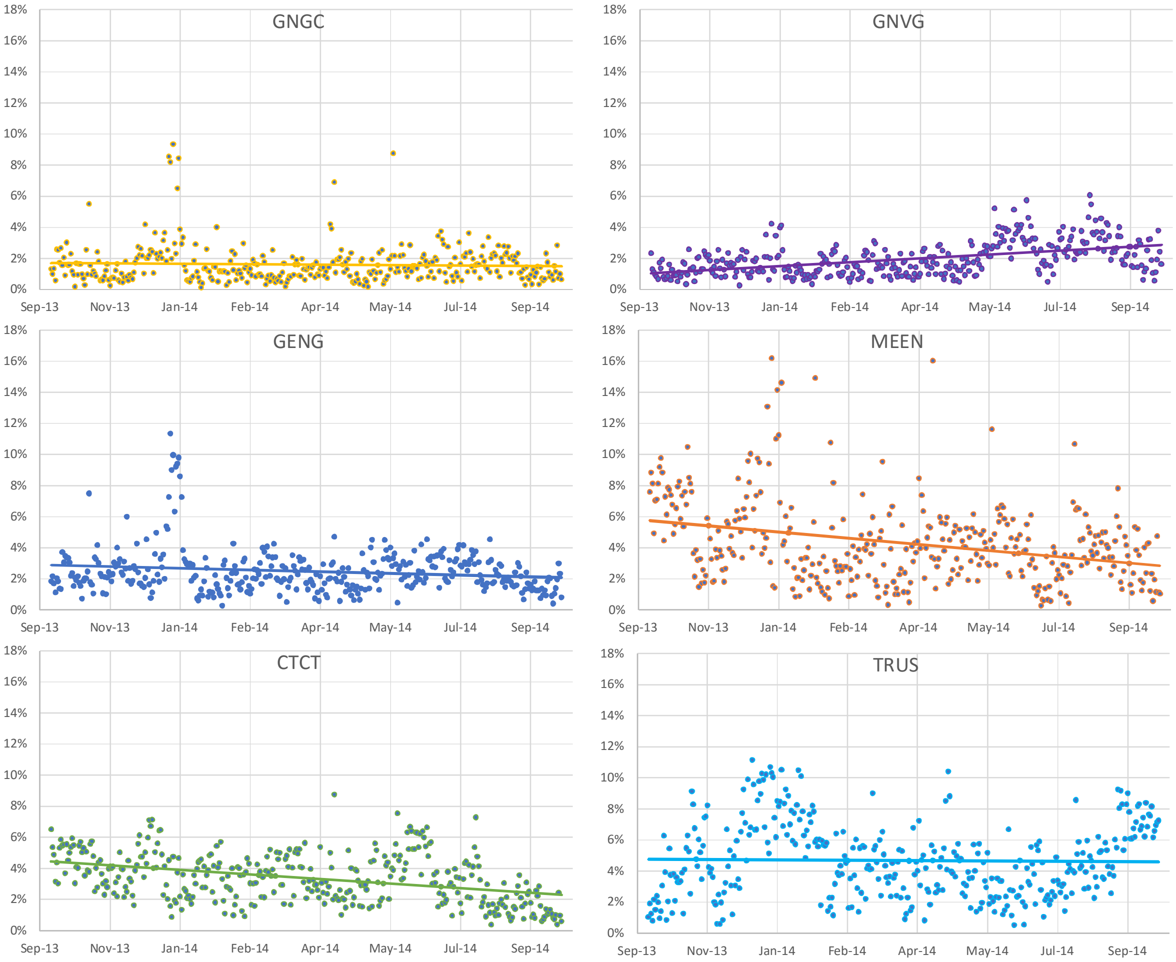

Figure 20 shows the average retailer percentage error for each day of the analysis period. Allocation groups are again combined for this section.

Some trends over time are evident. The D+1 predictions for MEEN and CTCT appear to improve over the year studied, while the GNVG predictions appears to become worse.

TRUS D+1 predictions were particularly poor in December 2013 and January 2014. This was when TRUS consumption was increasing rapidly in the BoP and North pools.

In general, the D+1 process produces poor predictions for December. There are many days in December that are nominally business days, but which have consumption more typical of a non-business day. This makes them difficult to model. A similar effect can be seen for Easter at the start of April.

4.8 Results conclusions

For most retailers, the D+1 process produces overall allocations that have a similar error to the current process.

Group 2 consumption is not modelled as well in the D+1 process as it is in the current process. This is unsurprising, since metered values are currently used for the Initial allocation. If it were desired, the total error for group 2 ICPs could be halved by moving the 50 most variable ICPs to group 1.

For most retailers, the D+1 process produces predictions for non-TOU consumption that are slightly worse on average than the current process. However, they often have a lower incidence of large errors, and this may be preferable for the purposes of pipeline management.

In practice, it is not important which allocation group uses a retailer's gas. Retailers are mostly concerned with balancing their nominations at a pool level. Conveniently, the D+1 predictions are more accurate when allocation group is ignored because of off-setting inherent in the modelling process.

GNGC is the only retailer that would experience significantly worse accuracy under a D+1 allocation. They have a large proportion of group 2 demand, and their non-TOU demand is more difficult to model. However, this mostly reflects that their current allocation error is very low. Under a D+1 allocation process, their absolute error would be similar to other major retailers.

Appendix A. Pipeline management

Retailers can use the information provided by D+1 allocation to assist with their nominations throughout the month. This will help to improve their nomination accuracy which will improve pipeline management.

If Market Based Balancing (MBB) is implemented, ROI will be cashed out each day that tolerances are exceeded. D+1 allocations will allow retailers/shippers more information to minimize their ROI, which should minimize cashouts.

One way in which shippers could use D+1 allocations to assist with nominations is shown in Figure 21 below.

Figure 21 - Example ROI for a retailer using D+1 allocations

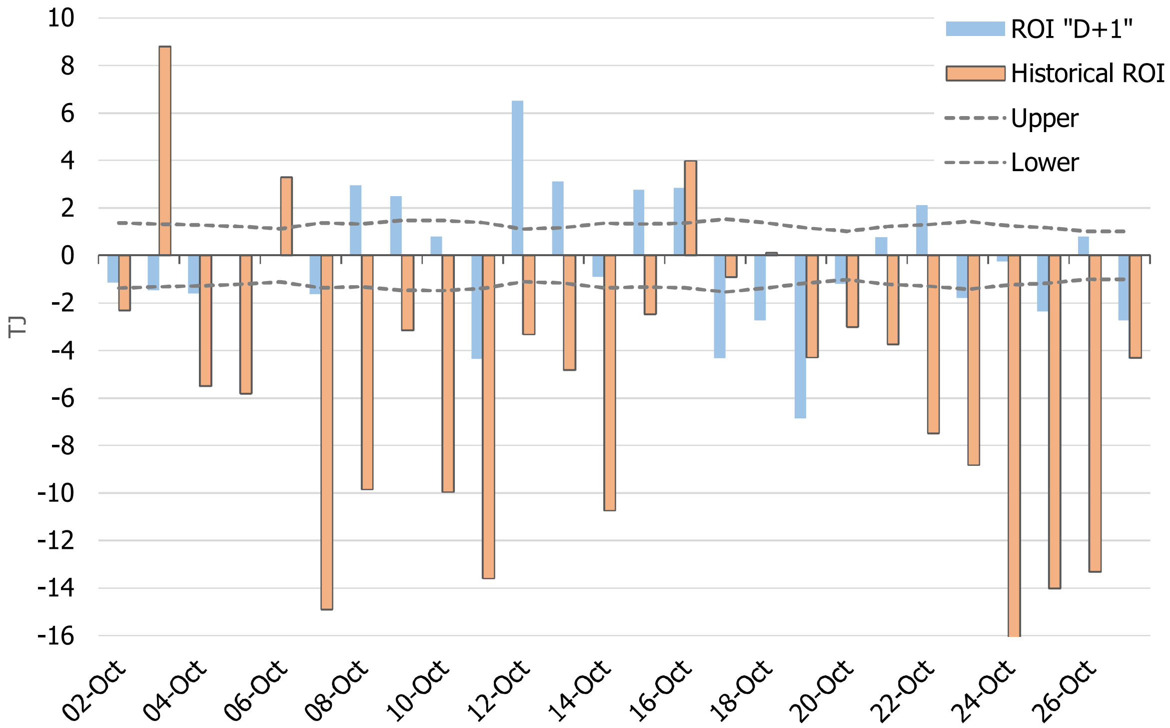

The nomination for a day is calculated by taking the previous day's allocation (the D+1 prediction), plus/minus any difference between yesterday's nomination and allocation.

In this simple example, cashouts occur on roughly half the days, and the ROI is much better than it was historically.

This is a very simple approach, and retailers could likely do much better with some effort[11].

[11] For example, considering the weather would improve the accuracy of forecasts.

Appendix B. Miscellaneous notes

4.9 Retailer groupings

In 2011 E-Gas collapsed and the majority of its customers were acquired by Nova. To account for this all historical E-Gas consumption was classified as Nova.

4.10 EDNZ SKF

EDNZ in the SKF pool had no group 6 ICPs for most of 2009, and these customer were instead classified as group 4. Because of this, the SKF pool prior to December 2009 has not been included in the development of models.

4.11 Other Groups

In the GAS050 data file there are some group 0 and group 3 ICPs. In many cases these were group 2 ICPs that have been moved to group 4, or ICPs that have been discontinued for some reason. These ICPs are dealt with by treating them as if they were in group 1. This means that they do not influence the allocation process, because their consumption is assumed to be known.

Appendix C. Retailer errors by pool

Figure 22 - Distribution of absolute errors for SKF pool for all consumption

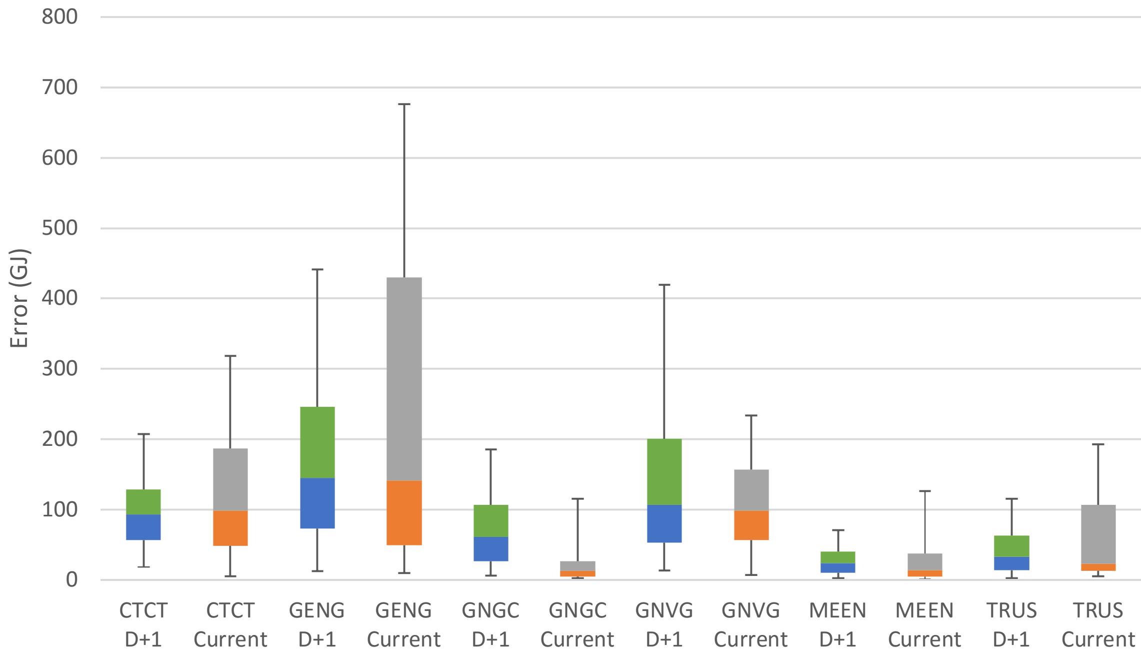

Figure 23 - Distribution of absolute errors for North pool for all consumption

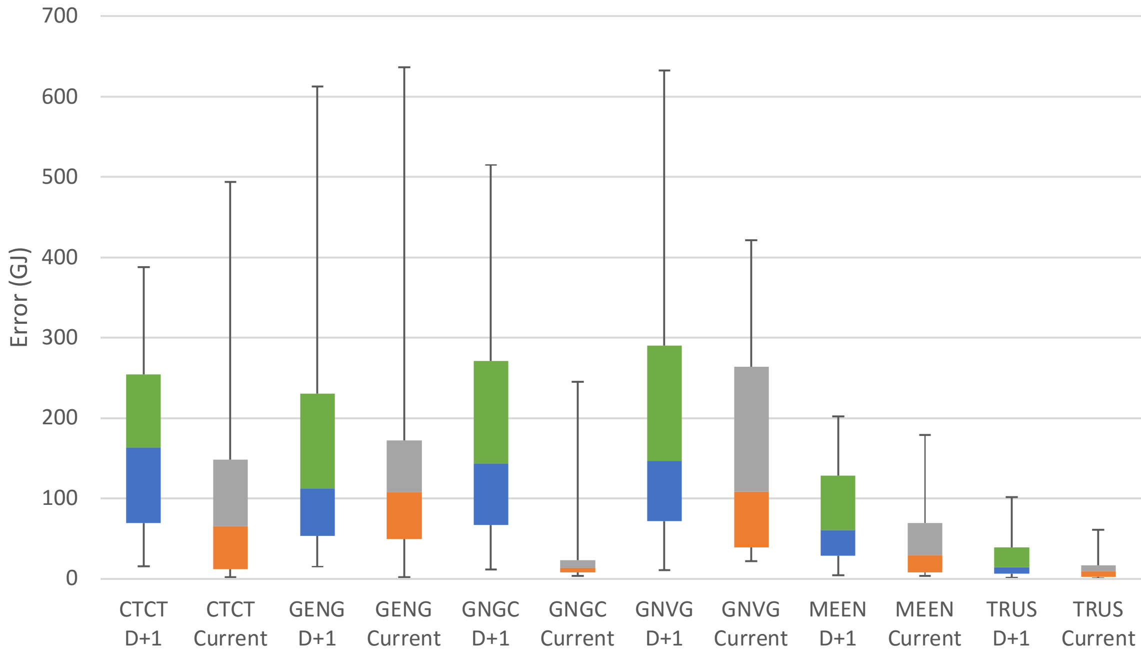

Figure 24 - Distribution of absolute errors for BoP for all consumption