ACIL ALLEN
CONSULTING

Briefing Note

Date: 18 August 2015

To: Rainer Korte
Simon Appleby
Bill Jackson

From: Jeremy Tustin
Timothy Weterings

Subject: Analysis of Market Impact Component of STPIS

ACIL ALLEN
CONSULTING PTY LTD
ABN 68 102 652 148
LEVEL NINE
60 COLLINS STREET
MELBOURNE VIC 3000
AUSTRALIA
T+61 3 8650 6000
F+61 3 9654 6363
ACILALLEN.COM.AU

BRISBANE
CANBERRA
SYDNEY
MELBOURNE
PERTH
ADELAIDE

This briefing note is to summarise ACIL Allen's analysis of the Market Impact Component (MIC) of the Australian Energy Regulator's (AER) proposed Service Target Performance Incentive Scheme Version 5 (TxSTPIS5).

On Tuesday, 11 August 2015 you asked us to analyse the MIC. In particular, you asked whether the MIC as proposed for TxSTPIS5 is systematically biased to give TNSPs a negative payoff.

The first sections of this briefing set out background, describing principles applicable to TxSTPIS5, the way the MIC will operate and the number of constrained dispatch intervals that occur in a given period.

Following that we outline some conclusions. In summary, our view is that the MIC as currently proposed has several statistical biases (i.e. more likely to give either a positive or negative payoff than the other):

— There is a bias to give an increasingly negative expected payoff to an increasingly efficient TNSP
— There is an indeterminate bias when the distribution of the number of dispatch intervals (DIs) in which a constraint binds is positively skewed
— There is a penalty for volatility in yearly constrained DI counts.

Overall, these factors create a negative bias in the payoff to a TNSP over time, such as over a regulatory period, though the bias in a single year is positive. None of these biases are consistent with the appropriate objectives of the scheme.

Background

Principles applicable to MIC payoff

The purpose of TxSTPIS5 and other incentive mechanisms for Transmission Network Service Providers (TNSPs) is to provide them with an incentive to improve the service they provide towards the efficient level of service (the minimum number of constrained DIs a TNSP can reasonably be expected to experience).

The MIC relates to the number of DIs in which a transmission constraint prevents AEMO from dispatching at the lowest cost from the available bids (constrained DIs).

In principle, the MIC should give a TNSP:

— a reward (positive payment) if it reduces the number of counts (incentive to improve)
— a penalty (negative payment) it if increases the number of counts (incentive not to deteriorate)
— no payment if the number of constrained DIs is unchanged.

If a TNSP maintains performance at the same level over time but this is above (worse than) the efficient level, the appropriate payoff is zero, although one could argue that this should be negative, to motivate improvement.

The number of constrained DIs is stochastic

In any given period the number of constrained DIs that will occur is subject to a large number of factors. Many of these (such as demand, generator bidding patterns, wind farm output) are beyond the TNSPs control.

In other words, there is a stochastic (random) element to the number of constrained DIs that will occur on a given network in a given period.

Therefore, the number of constrained DIs that occurs on a given network in a given period can be thought of as a random draw from a probability distribution. This means that actual number of constrained DIs observed in any given period will vary around an expected number of constrained DIs (mean) associated with the distribution.

In the long run, the mean of the constrained DI distribution will be centred around the 'target' for a network which is maintaining its current performance level. An 'improving' network is synonymous with a decreasing mean, although in any given year a larger constrained DI count might be observed.

The true distribution of constrained DI counts is not known and will not necessarily be the same for different TNSPs. However, the following properties are known, or can reasonably be assumed.

The number of constrained DIs that will occur in a given year has two components:

- relatively frequent small events cause a small number of constrained DIs each event (the background noise component)
- occasional events cause a large number of constrained DIs per event (the serious event component).
- Further:
- it is not possible for the number of constrained DIs to fall below zero
- it is not efficient for the number of constrained DIs to fall below the efficient level, because the cost of doing so outweighs the benefits. It is also increasingly unlikely that DI counts near and beyond 'E' will occur. Therefore, the distribution is likely to be truncated at or near the efficient level.[1]
- there is no (relevant) upper limit on the number of constrained DIs that can be observed in a given period on a given network.
- In lieu of a long time series of constrained DI counts, it may be reasonable to assume that the background noise component of the distribution is normally distributed.
- However, depending on the preponderance of serious events the distribution of constrained DIs could be viewed as positively skewed. That is, the median of the distribution is less than the mean and the positive tail is 'longer' than the negative tail.

This would look something similar to Figure 1.[2]

[1] For modelling purposes we have truncated it at this level.

[2] This is only illustrative. In this case the population mean is approximately 1000 and the median is approximately 607 (the efficient level is not depicted here).

FIGURE 1 – **ILLUSTRATIVE POSITIVELY SKEWED DISTRIBUTION**

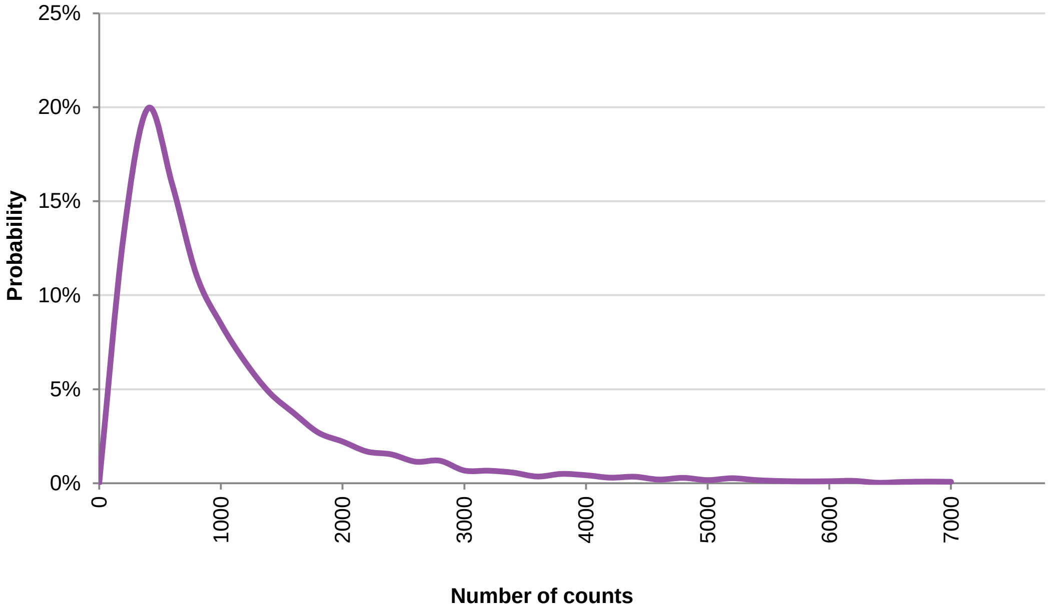

SOURCE:ACIL ALLEN CONSULTING

The MIC in TxSTPIS5

The MIC as proposed in TxSTPIS5 consists of:

— a target 'T' – set at the rolling three year average number of constrained DIs

— a performance measure 'P' – set at the rolling two year average number of constrained DIs

— a payoff range - ± 1% of Maximum Allowable Revenue

— two target range parameters which specify the 'distance' from the target 'T' over which a reward or penalty is received, α and β. In the proposed TxSTPIS5 these parameters are not mentioned but both are set to 1. They are also not described as different than one another, but conceptually this is possible

— a cap – no single event can contribute to total counts by more than 17 per cent of the target.

— a minimum target – T cannot be set below 100.

In a given period the number of constrained DIs is drawn from the TNSP's distribution and used, with the number for the previous year to calculate 'P'. The TNSP's payoff is:

1. if P=T-αT, payoff is +1% of MAR[3]
2. if T-αT < P < T, payoff is between 0 and +1% of MAR
3. if P=T, payoff is zero
4. if T < P < T+βT, payoff is between 0 and +1% of MAR
5. if P=T+ βT, payoff -1% of MAR . [4]

Figure 2 illustrates.

[3] Currently T-αT is 0
[4] Currently T+ βT is 2T

FIGURE 2 – **MIC PAYOFF RANGES**

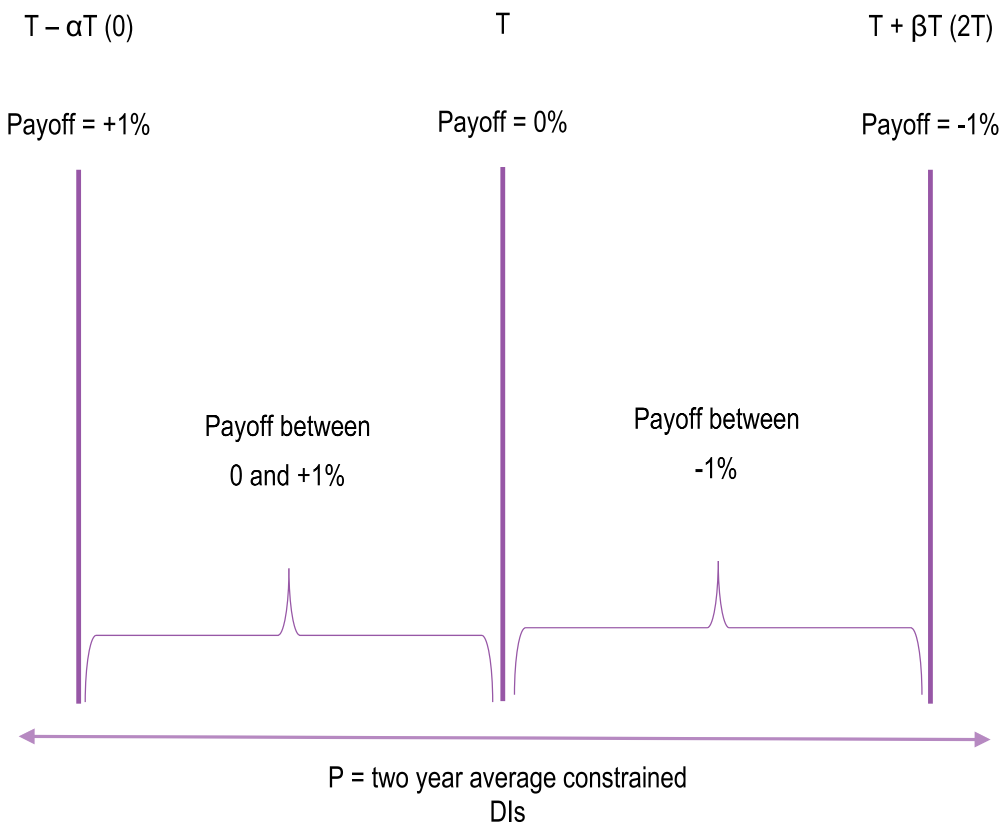

SOURCE:ACIL ALLEN CONSULTING

The expected MIC payoff for a TNSP

The expected MIC payoff for a TNSP is calculated by multiplying the probability of receiving any given payoff with the payoff itself. Payoffs are continuous, so the mathematics are complex. Conceptually, it can be described by reference to Figure 3 as follows:

In this illustration the efficient number of constrained DIs is set (arbitrarily) to 20% of the mean, so the maximum payoff the TNSP can receive in any given period is 0.8% of MAR.

In summary, the payoff is the sum of:

— Purple area multiplied by payoffs between 0 and +0.8%[5]
— Grey area multiplied by payoffs between 0 and -1%
— Black area multiplied by -1%.

Whether the expected payoff is positive, negative or zero depends on the relative sizes of the purple, grey and black areas in the TNSPs actual probability distribution.

[5] The particular value, 0.8%, is the result of setting E to 80% of T. The level is unimportant. If E is greater than 0 then the upper bound of the payoff is less than +1% of MAR.

FIGURE 3 – **ILLUSTRATIVE MIC PAYOFF**

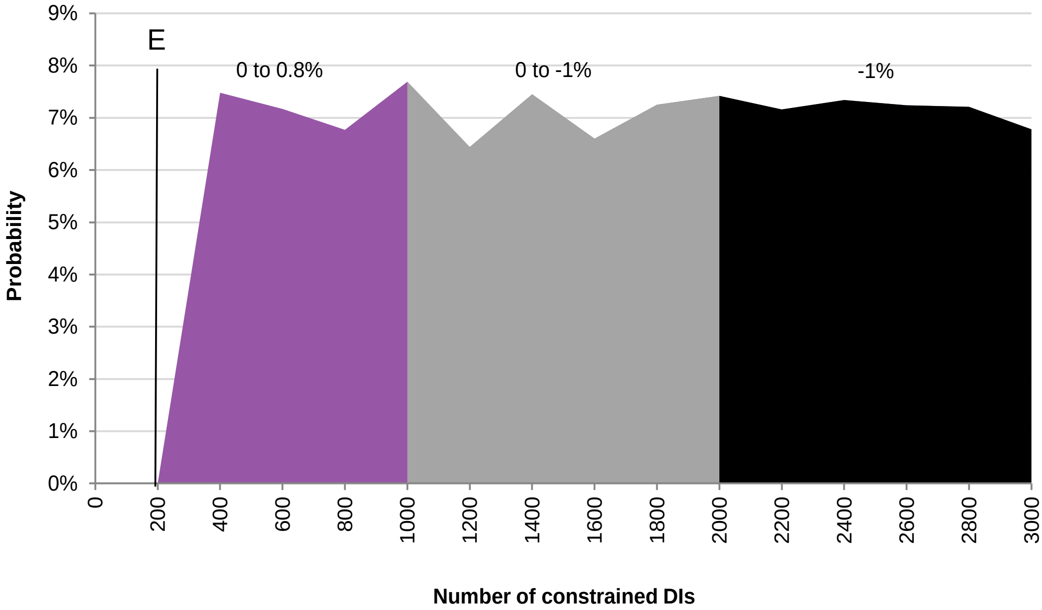

SOURCE: ACIL ALLEN CONSULTING

Results

This section sets out the results of our analysis of the MIC.

The MIC as proposed gives a negative expected payoff to an efficient TNSP

We define an efficient TNSP as on for which the expected value of P is the efficient level of constrained DIs. [6]

It is neither desirable nor likely that a TNSP will experience less than the efficient number of constrained DIs. Therefore, a disproportionate number of the stochastic outcomes are associated with negative payoffs. In a sense, this TNSP can experience bad luck much more readily than good.

Figure 4 shows the payoff for a TNSP which is almost efficient. In this case we assume a normal distribution for constrained DIs, but the particular shape is unimportant. Outcomes in the yellow area are part of the distribution, but are impossible, because this is below the efficient level.

For this TNSP, the expected payoff is the sumproduct of:

— the purple area, with payoffs between 0 and 0.8% of MAR
— the grey area and payoffs between 0 and -1% of MAR
— the black area and -1% of MAR.

Simply put, the sum of the area of the grey and black areas is larger than the purple, so the expected payoff is negative. [7]

If the outcomes in yellow are instead assumed to be at the efficient level (i.e. stacked up on ‘E’) the payoff is less negative.

[6] The expected value of P is the mean of the distribution of the number of constrained DIs
[7] This is not a formal description, the shape of the purple and grey areas is relevant to the payoff as well.

FIGURE 4 – **ILLUSTRATING EXPECTED MIC PAYOFF**

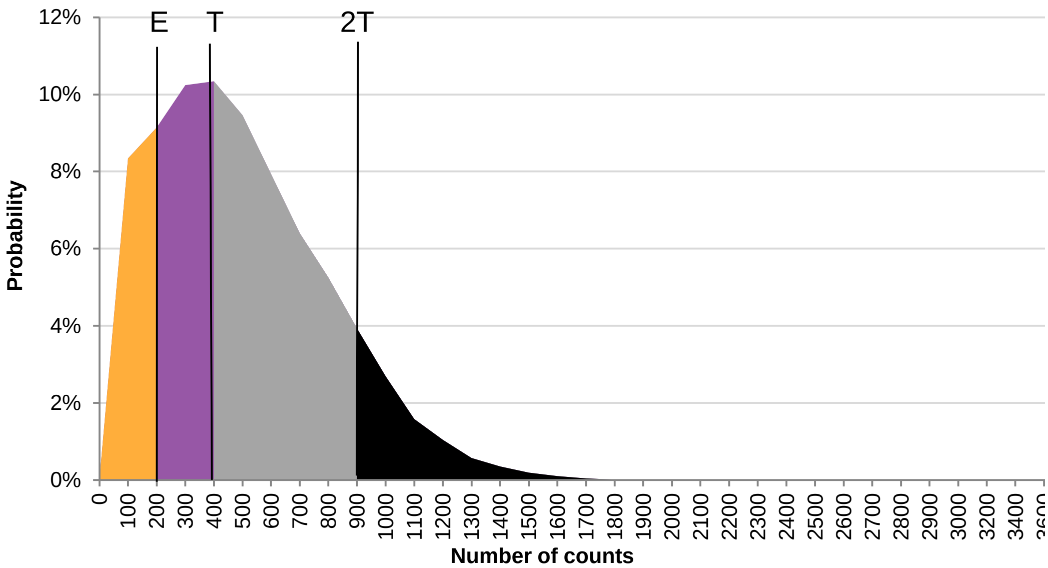

SOURCE: ACIL ALEN CONSULTING

The existence and size of a negative bias depends on the ratio of 'E' to P (and T). In simulations with normally distributed number of constrained DIs with declining means the results were as shown in Figure 5. As the TNSP becomes more efficient, and thus P and T decline towards E, the expected payoff becomes more negative.

FIGURE 5 – **NEGATIVE MIC BIAS WITH MINIMUM EFFICIENT NUMBER OF CONSTRAINED DIs**

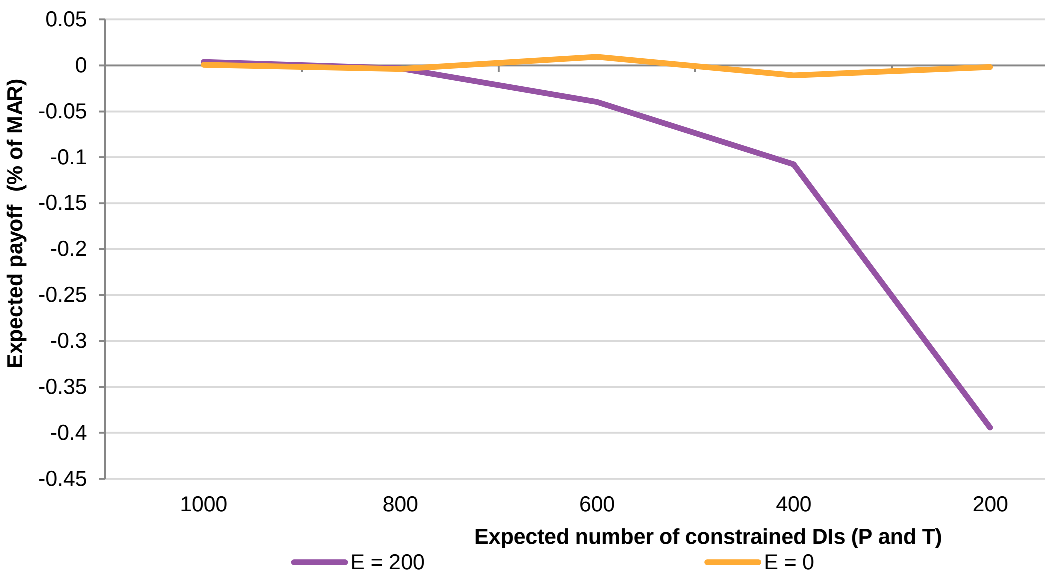

SOURCE: ACIL ALLEN CONSULTING

Expected payoff is biased for skewed distribution of P

When a positively skewed distribution is used the MIC appears to have a positive bias in a given year.

Figure 6 shows the payoff ranges from Figure 4 with a positively skewed distribution overlaid (in this case a log-normal distribution). The probability that the TNSP will experience a payoff in each range is the area under the curve in that range.

In this case, E is again set arbitrarily to 20% of the expected value of P, so the maximum payoff the TNSP can receive in any given period is 0.8% of MAR.

As above, the payoff is the sum of:

— purple area multiplied by payoffs between 0 and +0.8%
— grey area multiplied by payoffs between 0 and -1%
— black area multiplied by -1%.

FIGURE 6 – **MIC PAYOFFS WITH SKEWED DISTRIBUTIONS**

Payoff = 0.19% MAR

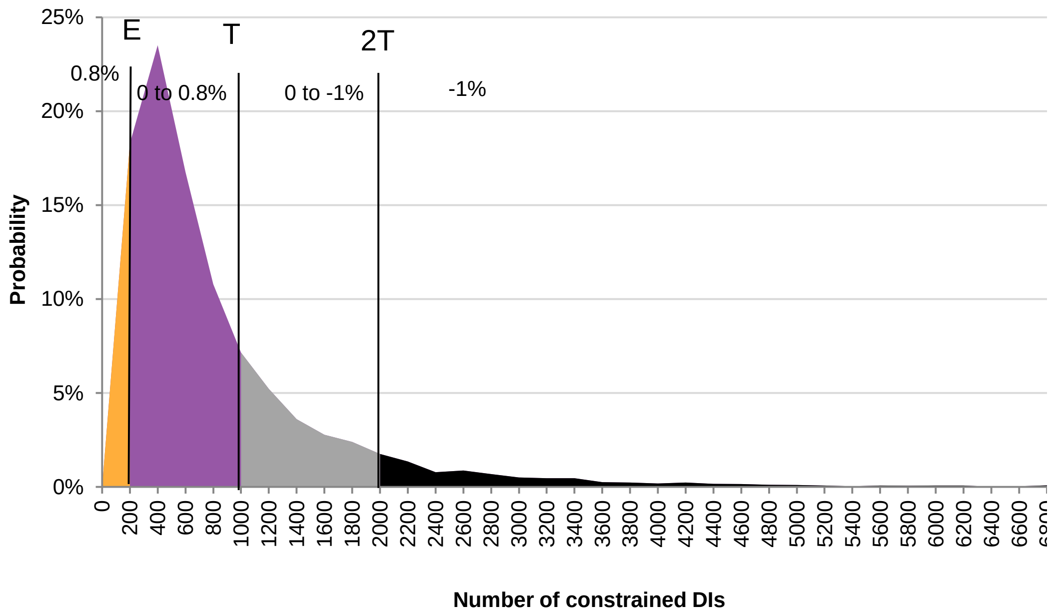

Payoff = 0.56% MAR

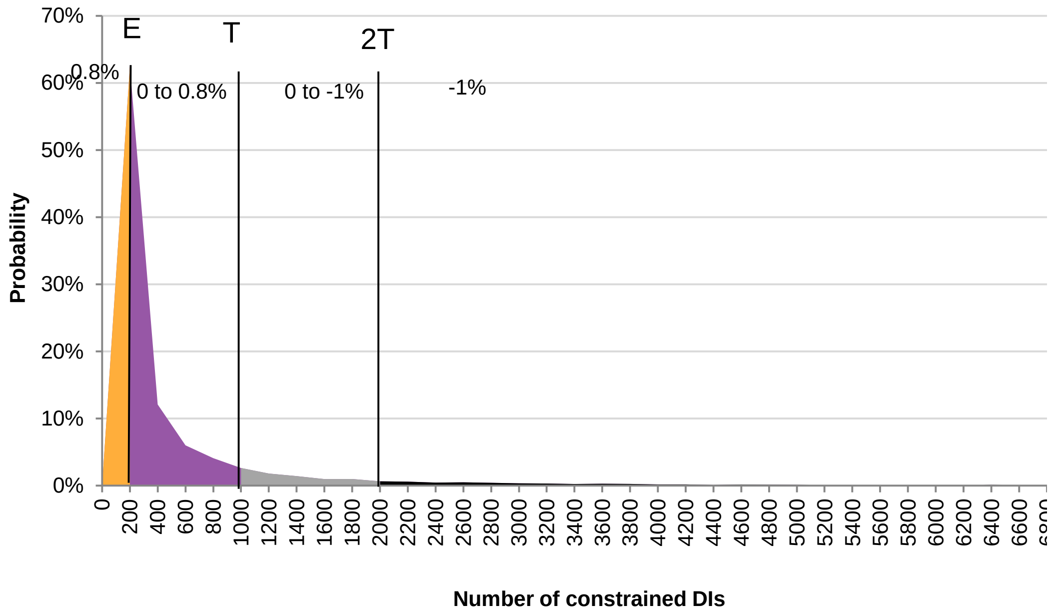

SOURCE: ACIL ALLEN CONSULTING

The top pane of Figure 6 shows a modestly skewed distribution (lognormal distribution with mean=1000, Standard deviation = 1). In this case the payoff is biased in a positive direction by approximately 0.18 percentage points.

The second pane shows a less skewed distribution. There is a very small chance of a very large 'P' (including the effect of the 17% cap) and 'good' years are substantially more likely than 'bad'. In this case the bias is substantially positive, approximately 0.56 percentage points.

Dynamic impact – the MIC is negatively biased over time

Our third conclusion is that the MIC payoff is negatively biased over time. The reason is that the value of a constrained DI is a function of the target so it is not constant over time.

Consider a TNSP's average payoff over the 5 years calculated using the AER's model. The model shows, in cell AD32, the total (sum) payoff for years 7 to 11 (inclusive) so the analysis is applied in those years. Year 7 can be thought of as the first year of the regulatory period. , To see the negative bias, first consider a TNSP that experiences constrained DIs as shown in Table 1. In this case the TNSP 'sits' at a level of 1,000 constrained DIs each year. P and T are at this level and its payoff is zero over the ten years shown.

TABLE 1 – **PAYOFF EXAMPLE FOR TNSP WITH CONSTANT CONSTRAINED DIS**

Year	Constrained DIs	P	T	Payoff
1	1000	1000	1000	0.000
2	1000	1000	1000	0.000
3	1000	1000	1000	0.000
4	1000	1000	1000	0.000
5	1000	1000	1000	0.000
6	1000	1000	1000	0.000
7	1000	1000	1000	0.000
8	1000	1000	1000	0.000
9	1000	1000	1000	0.000
10	1000	1000	1000	0.000
Total	**10,000**	**10,000**	**10,000**	**0**

SOURCE: ACIL ALLEN CONSULTING

Now consider Table 2, in which the TNSP has one bad year (in year 2) and 'makes good' in later years (3 and 4). The total number of constrained DIs over the period shown is the same. However, due to changes in the target and the value of a DI In different years, the TNSP receives a net penalty over the period shown.

The result is the same if the bad year occurs after the 'good' (Table 3) and occurs regardless of the number of 'good' and 'bad' years. If the '*overs and unders'* sum to zero, the payoff is negative over time. The reason is that the value of DIs Is linear in a year, but not constant between years. The smaller the target (T), the more valuable (and costly) a single constrained DI. The larger T, the less valuable is one constrained DI.

TABLE 2 – **PAYOFF EXAMPLE FOR TNSP WITH CONSTANT CONSTRAINED DIS**

Year	Constrained DIs	P	T	Payoff
1	1000	1000	1000	0.000
2	1500	1250	1000	-0.250
3	750	1125	1167	0.036
4	750	750	1083	0.308
5	1000	875	1000	0.125
6	1000	1000	833	-0.200
7	1000	1000	917	-0.091
8	1000	1000	1000	0.000
9	1000	1000	1000	0.000
10	1000	1000	1000	0.000
Total	**10,000**	**10,000**	**10,000**	**-0.073**

SOURCE: ACIL ALLEN CONSULTING

TABLE 3 – **PAYOFF EXAMPLE FOR TNSP WITH CONSTANT CONSTRAINED DIS**

Year	Constrained DIs	P	T	Payoff
1	1000	1000	1000	0.000
2	750	875	1000	0.125
3	750	750	917	0.182
4	1000	875	833	-0.050
5	750	875	833	-0.050
6	1000	875	833	-0.050
7	1000	1000	917	-0.091
8	1000	1000	917	-0.091
9	1000	1000	1000	0.000
10	1000	1000	1000	0.000
Total	**10,000**	**10,000**	**10,000**	**-0.025**

SOURCE: ACIL ALLEN CONSULTING

Bringing the components together

Our final conclusion is that, considering the conflicting biases considered above, the MIC as currently proposed is biased to provide a negative payoff to a TNSP over time.

We simulated the average return to a TNSP over a five year (regulatory) period using the distributions of constrained DI's considered previously. We did this 10,000 times to get a distribution of regulatory period returns.

Figure 7 shows this distribution of average returns over a regulatory period.

FIGURE 7 – **DISTRIBUTION OF AVERAGE RETURNS OVER A REGULATORY PERIOD UNDER A LOGNORMAL CONSTAINED DI DISTRIBUTION**

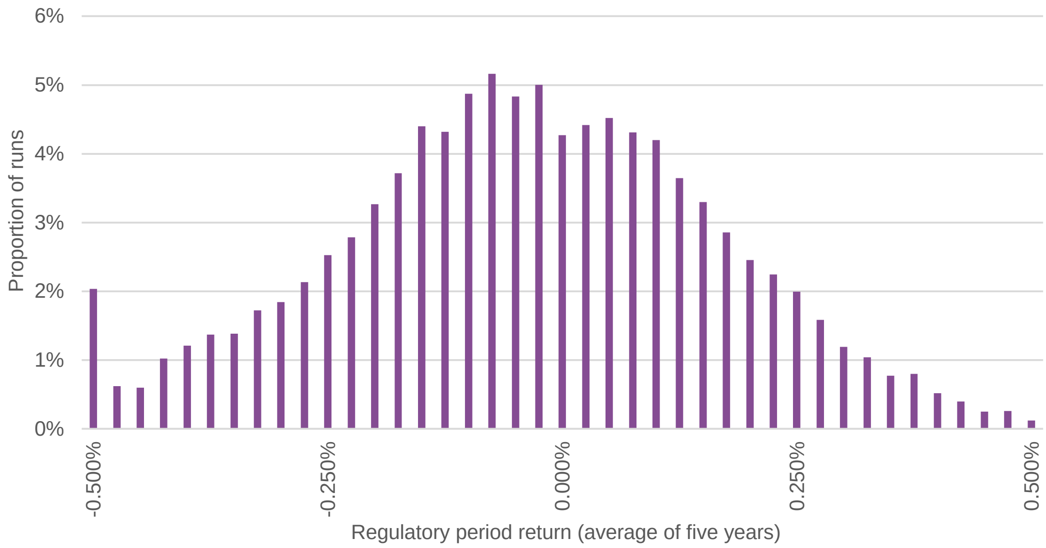

[a] this is a reference note

Note:

SOURCE:

On average, the regulatory period (five year) payoff was -0.049 per cent. 95 per cent of the time the return was between -0.456 and 0.351. This volatility reflects the variability in initial conditions and the lagged nature of the target calculation. This shows that in any given regulatory period, the payoff for a TNSP which is neither improving nor worsening, could be positive or negative. However, the issue of bias is concerned with the average of this distribution.

The *central limit theorem* can be applied to determine whether the regulatory period return is biased, using a hypothesis test. If we define the regulatory return as X, then the hypotheses are:

H0: The average regulatory period return, μ, is equal to zero

Ha: The average regulatory period return, μ, is not equal to zero (i.e. it is biased)

We know that $\bar{X}$ = -0.049. The standard deviation of X is 0.204. This information is used to calculate a t-statistic of -22.1.

If μ is in fact zero, 95 per cent of the time the t-statistic will be between -1.96 and 1.96. As our observed t-statistic is less than -1.96, we reject the null hypothesis and conclude that the average regulatory period return is not equal to zero, at the 95 per cent confidence level (i.e. returns are biased).

This test was also conducted using normally distributed constrained DIs, and the same conclusion was drawn. In this case the average regulatory return was -0.032.

Conclusion

In summary, our conclusion is that a TNSP with:

1. a lognormal distribution of constrained DIs (which is positively skewed)
2. T set at the mean of that distribution initially

will expect to achieve a positive return in a given year, even if their underlying mean number of constrained DIs does not change (i.e., without systematic improvement).

Notwithstanding the positive bias in the annual return, over time, whether a regulatory period or longer time, their expected payoff from the proposed MIC is negative.

The negative bias results from the way 'T' follows the TNSPs performance and the way that the value of a single DI varies from year to year when it does. The latter is due to the 0 <T < 2T structure of the payoff ranges.

The result is that the proposed MIC punishes TNSPs for the volatility that is inherent a stochastically determined value such as constrained DIs. This is not consistent with the objectives of the scheme.