

Cessnock Council Efficiency Report



Disclaimer

This Report was prepared by Professor Drew, Professor Miyazaki and Professor Ferreira on behalf of Professor Joseph Drew Pty Ltd. The work herein has also been independently assured by Professor Yunji Kim from Seoul National University. This Report was produced in a strictly independent manner and the opinions expressed in the Report are thus exclusively the views of its authors and do not necessarily coincide with the views of the local government or any other body. The information provided in this Report may be accurately reproduced for media review, quotation in literature, or non-commercial purposes, subject to the inclusion of acknowledgement of the source and provided no commercial use or sale of the material occurs.

The information contained herein is believed to be reliable and accurate. However, no guarantee is given as to its accuracy or reliability, and no responsibility or liability for any information, opinions, or commentary contained herein, or for any consequences of its use, will be accepted by Professor Joseph Drew Pty Ltd, or by any person involved in the preparation of this report.

Cessnock Council Efficiency Report

Executive Summary

This report examines a number of efficiency metrics spanning crude ratios all the way up to sophisticated envelopment analyses. We also briefly consider the determinants of efficiency with respect to the characteristics of Cessnock. The picture which emerges is a local government with commendable cost control, although actual conversion of inputs into outputs is less flattering (but still good in view of the unique operating environment faced). Towards the end of this report, we list a number of measures that might be expected to improve various aspects of efficiency. Furthermore, in the appendix we list the efficiencies proposed by Cessnock City Council, with our assurance or comment as appropriate. In sum, there is only marginal improvements to efficiency that could be reasonably expected of Cessnock City Council and we do not believe that these will have any material impact on financial sustainability over the medium-term. We certainly believe that all of the plausible efficiencies contained in this report should be pursued with the utmost vigour, but they will sadly not be a substitute for a special rate variation (SV).

1. Introduction

It appears that both the Office of Local Government (OLG) and the Independent Pricing and Regulatory Tribunal (IPART) are concerned about the efficiency of local governments seeking to increase their rates above the prescribed cap.

Unfortunately, it also seems that there is a good deal of confusion about what precisely efficiency is, how to measure efficiency competently, as well as the potential for efficiency improvements to put material downward pressure on taxation.

Efficiency is often ill-defined in a public policy sense despite the fact that economists have quite precise definitions and ways of measuring it. Typically, scholars make reference to three distinct kinds of efficiency which local governments exert varying levels of control over: (i) allocative, (ii) dynamic, and (iii) technical efficiencies.

Allocative efficiency refers to how scarce resources are harnessed to enhance the flourishing of citizens (Fergusson, 1972). To achieve allocative efficiency, it is necessary for decision-makers to carefully direct inputs to both the quantity and quality of goods and services desired by the community. In a local government sense the principal mechanism for allocative efficiency is the democratic process over time.

Dynamic efficiency, by way of contrast, refers to changes to allocative or productive efficiency over time (Drew, 2021). Dynamic efficiency is principally driven by improvements to learning or technology. Dynamic efficiency might also alter due to changes in regulatory practice or alterations to legislation, albeit typically in a deleterious manner. Dynamic efficiency largely arises due to the actions of others

(advancement in educational offerings or industrial products) and lies beyond the direct control of Councillors or local government management.

The third type of efficiency is technical (also referred to by scholars as productive or sometimes x-efficiency) and this refers to the optimal conversion of inputs into a large range of local government outputs (Drew, 2021). The inputs to the production process are staff and money and the outputs are too numerous to list (hence economists typically use proxies for the main types of goods and services produced by local governments). The state government, regulators and some citizens have put considerable emphasis on the concept of technical efficiency presumably believing that: (i) efficiency is a legitimate goal of government, and (ii) that efficiency might ineluctably lead to improved sustainability and/or lower taxes.

There is no good reason to think that efficiency is either a legitimate goal of government, or indeed that high levels of efficiency are even possible (Drew, Razin and Andrews, 2018). Scholarly work on public values has identified that citizens care most strongly about notions such as access to services, privacy, equity, civil rights, as well as safety and security (see, for example Bozeman, 2019). Efficiency rarely rates a mention unless citizens are confronted with a request to pay the full price for the services that they consume (Drew, 2021). Indeed, many of the things that citizens expect their governments to do are completely contrary to efficiency – for instance disaster response (whereby governments often have to pay penalty rates and the like to ensure quick relief for those suffering) or holding regular elections (considerable resources are expended for no additional goods or services output). We doubt very much that citizens would ordinarily argue that government functions such as these ought to be sacrificed in the name of efficiency. Moreover, it has long been held by scholars that efficient delivery of goods and services is inconsistent with democratic government in any case (see, Fenwick, 1920; Friedman, 1993). Indeed, we have only to briefly consider the disaster wrought at the hands of new public management¹ proponents to understand the folly of myopically pursuing efficiency in a democracy (see, for example, O'Flynn, 2007; Drew, 2021).

Nevertheless, regulators have continued to place strong focus on technical efficiency, in particular. The assumption seems to be that improvements to efficiency will result in higher sustainability or lower taxes. However, the scholarly evidence on this matter does not support this assumption (see Drew, Kortt and Dollery, 2015a). The main reason for this lack of support is that efficiency is a short-run concept, whereas sustainability (and tax rates in the context of a rate cap regime) are long-run matters. Any marginal changes to efficiency in the present are thus likely to pale into insignificance when set against decisions taken over many decades regarding the construction of infrastructure, addition of services, drawing down of debt, or the neglect to charge an average tax price for a local government area (the cumulative effect of this last factor is certainly a large contributor to Cessnock's predicament as demonstrated in our *Capacity to Pay* report). Indeed, one only has to consider the

¹ This was a public policy theory that tried to make government operate like business, in pursuit of efficiency. It was characterised by Bevan and Hood (2006) as 'targets and terror' and was largely an abject failure – not least because government is demonstrably not a business (Drew, 2025).

personal budget metaphor to understand the fallacious nature of pervasive assumptions in this area².

In a local government sense, the way to improve technical efficiency is to combine the optimal mix of production factors to produce a given quantity of outputs (what is referred to as an input-orientation). This is the role of local government managers. Presumably this is the focus of regulators, although as we shall see, their crude ratios are entirely incapable of measuring technical efficiency.

In the past regulators have sought to measure technical efficiency through a crude ratio defined as operational expenditure divided by population. In 2015 it was asserted that to be efficient a linear trend would need to be downwards sloping over a five-year³ period. This approach entirely neglected to consider how different factors of production might be best combined, and eschewed the time value of money altogether. Furthermore, the 2015 attempt at measuring efficiency also used the incorrect functional unit – it has been shown countless of times that in Australian local government, that number of properties is a superior denominator in the absence of more sophisticated weighted methods (Drew and Dollery, 2014; indeed, road lengths – the single largest item of expenditure – are negatively correlated to population size!). Moreover, a number of other serious problems exist with ratio approaches that we shall enumerate later. In sum, the crude metric still used in NSW is fatally flawed and only likely to lead to erroneous conclusions. Clearly something more sophisticated is required to allow valid statistical reasoning to take place.

In this report the centrepiece of our work are: data envelopment analysis (DEA) and free disposability hull analysis (FDH). This is world's best practice and sophisticated empirical work conducted by one of the leading scholars in academia today (Professor Ferreira). It is the only way to competently appraise the efficiency of Cessnock over time, and we conduct these analyses over an eight-year panel for the entire cohort of relevant NSW local governments⁴.

The remainder of this report is set out as follows. In the next section we review a number of ratio metrics that will provide an overview of relative performance compared to councils that the OLG deem to be similar to Cessnock. We will also present the results of a regression analysis which can be used to further assess the actual expenditure on staff against what might typically be expected. Following this we present world's best practice sophisticated DEA and FDH analysis. Thereafter, we conduct a DEA and FDH of tax efficiency. We also search for the determinants of efficiency and briefly outline the Council's efficiency journey. The report concludes with some observations regarding the potential for efficiency improvements to

²If a person went on an efficiency drive, they might hope to shave off a few percent on discretionary expenditures (savings on non-discretionary items such as food and water are usually not possible). Marginal savings of this kind would take many years to have a material impact on debts taken out to purchase property or the like, and pale into insignificance when set against the pecuniary implications of past decisions relating to things such as one's choice of occupation, marriage, or child-raising.

³A linear trend was not appropriate for data which was not linear, and five years is generally not considered sufficiently lengthy to establish a trend of the kind envisaged.

⁴The tax efficiency work is only conducted over a seven year panel of data.

materially alter the required special rate variation which needs to be passed on to taxpayers.

2. Ratio Analysis of Efficiency

Typically, for Special Variations (SV) councils a few so-called efficiency ratios are compared to try to make an argument about their relative technical efficiency. As we have already foreshadowed, this approach is flawed and can thus lead to completely erroneous conclusions. Nevertheless, it seems somewhat *de rigueur*, and the exercise will at least highlight the importance of the sophisticated work that forms the centrepiece of this report.


For the ratio comparisons that follow, reference is made to the peer group that draws on the OLG preferred categorisation.

Table 1. Peers Used in Comparisons

Bathurst	Kempsey	Singleton
Dubbo	Lismore	Tamworth
Eurobodalla	Mid-Western	Wagga Wagga
Goulburn Mulwaree	Queanbeyan-Palerang	Wingecarribee
Griffith	Richmond Valley	

The most efficient way of comparing Council to the peer group is to chart a box and whisker plot. Figure 1 provides details regarding how to interpret these plots:

Figure 1. Interpreting Box and Whisker Plots



In Figure 2 we present the OLG preferred metric of operational expenditure per capita as used during the 2015/16 *Fit for the Future* debates. As we have already suggested, this metric is completely flawed, and distinguished scholars have previously pointed out that it ‘simply does not measure efficiency’ (Drew and Dollery, 2015, p. 86).


According to Figure 2, Cessnock has consistently recorded the absolute lowest operational expenditure per capita for each of the last four years, relative to the peer group. This is suggestive of exceptional efficiency, however it would be very unwise to make decisions based on this result alone. First, the ratio depends on known unreliable data – population figures in intercensal years are merely estimates which the Australian Bureau of Statistics (2022) themselves have declared to typically impute errors of up to 8.9 percent at the SA2 level (typically several SA2 units need to be combined to produce local government level data). To see the importance of this problem one simply needs to compare population figures produced by the ABS for a particular year, with the rebased (corrected) data that is quietly substituted a year or so after each census – the differences are stark (including population growth which subsequently becomes revised to be population decline and vice versa).

Second, the majority of services in the Australian local government milieu are still delivered to properties (Drew, 2021). Using population as the denominator implicitly asserts that the cost of providing services such as roads and street lighting to a household of, say, four people is somehow four times larger than the cost of providing services to a single person household. It also seems to suggest that if a new baby is born to a previously childless couple that the cost of providing services (such as the aforementioned roads and street lights) increased by fifty percent; but if the baby was born to a family that already had two children then council costs only increased by twenty-five percent. Clearly, the idea of population as the functional unit of efficiency fails even the most basic logic tests and is quite untenable.

Third, the metric also implies that the cost of providing services to people living on farms is somehow comparable to the cost of providing services to people living in town. It might even be construed to suggest that there is no cost for providing services to business – especially non-retail establishments uncorrelated with population size. Clearly, these implicit assumptions cannot be true, especially given that the various different categories of properties do not receive anything like the same basket of local goods and services.


Fourth, operational expenditure per capita ignores the single largest item of costs for local government in NSW – roads. Indeed, roads are negatively correlated to population size ($r = -0.2531$) – this fact also further confirms that the output from this ‘efficiency’ ratio is likely to be quite misleading.

Figure 2. Operational Expenditure per Capita



In Victoria operation expenditure per property assessment is used instead. In Figure 3 we present the metric for Cessnock council relative to the peer group. In this instance, Cessnock is not only the lowest spending council in each year, but also an extreme outlier. However, operational expenditure per assessment, whilst better than the NSW metric, is still flawed – importantly, use of a metric of this kind still means that we must (implausibly) subscribe to the assumption that all categories of ratepayers receive more or less equivalent services. One merely needs to drive around greater Cessnock to understand that this assumption is not true. The metric used in Victoria also continues to ignore the single largest item of expenditure – roads. {only if we believed that all types of properties had similar length and types of road frontage could this neglect be tolerated}. For all these reasons, it would be unwise to place undue reliance on the results presented in Figure 3 either.

Figure 3. Operational Expenditure per Property Assessment (\$)




As we have already stated, the only way to competently assess efficiency is to use a sophisticated empirical technique that is capable of measuring the conversion of the various production inputs into multiple and appropriate proxies for outputs (even this approach is not perfect because of the failure of NSW to include a consistent measure of service quality like is done in jurisdictions such as Victoria). Before doing so, in the next section, we will examine a few further metrics that will provide some additional context for earlier discussions, and also expose misconceptions typically held by people in the community.

In Figure 4 we provide details of staff expenses per property assessment. It is important to do so, because inevitably in any SV at least one person will claim that all the problems stem from over-staffing or something of this kind. Indeed, we are aware that angry sentiments have already been expressed by a tiny minority of the community towards staff in social media – this is simply not acceptable: staff should not be exposed to psycho-social or physical risk because some people are unhappy with the facts confronting Cessnock City Council. We must be clear – there are no criteria that IPART can apply to reject or reduce an SV application because of angry words or threats; facts and constructive feedback are the only criteria supported by OLG Guidelines and the NSW Local Government Act (1993, NSW). We therefore urge restraint and constructive engagement by people who might be unhappy with the situation (we also again call on the NSW Government to make appropriate changes to the rate cap regime and associated SV process before serious harm is sustained by someone).

In Figure 4 below, we show that staff unit costs for the Council are the absolute lowest in the peer group for each of the four years. We note that there have been


increases of late in an absolute sense, but that the relative position remains the same. End users should be mindful that most staff (in any sector) receive pay increases annually in line with the relevant Award.

Figure 4. Staff Expenditure per Assessment (\$000)



We have noted earlier that technical efficiency is the conversion of inputs into outputs, thus it behoves us to also look at the proportion of expenditure at Cessnock on staff. In Figure 5 we do so and find that Cessnock has been in the top quartile (top twenty-five percent of peers) on two of the four occasions. This either means that council is outsourcing relatively more, or that materials and other expenses are also relatively lower than peers – or perhaps both. {This result is likely to be important when we turn to the more sophisticated envelopment analyses later in the report}. We will review the other major components of expenditure shortly, but what we have in front of us (Figure 5) certainly points to the need for more sophisticated analysis of efficiency such as the DEA and FDH which will ultimately follow.

Figure 5. Proportion of Expenditure on Staff (%)



Before leaving the matter of staffing we present the details of an econometric assessment of the expected level of staff expenditure for a council of Cessnock's characteristics. As we noted in the *Capacity to Pay Report*, econometrics is the sophisticated mathematics routinely performed by economists and many other scholars. This kind of work needs to be done by *bona fide* experts – and typically the training includes both undergraduate and postgraduate tertiary study. The report authors are attested to by hundreds of scholarly publications and thousands of citations by their peers; furthermore, the lead author of this report is an editor at a highly-ranked empirical journal – thus the estimates that follow are beyond reasonable doubt.

To produce the econometric predictions, we regressed staff expenditure against a suite of regressors long used by scholars to produce Australian local government cost functions (see, for example, Drew et al., 2021 for just one of countless examples of this kind of work; also read the *Capacity to Pay Report* for further details about econometrics). The regression had an extremely high coefficient of determination (0.9443) which means that the variables used nicely explained most of the typical staff costs.

In Table 2 we detail the predicted and actual staff costs for Cessnock for the last three years of the panel – and we consider this is to be important information for senior management at Cessnock because it provides a reliable guide regarding what could typically be expected for a council of this kind in NSW:

Table 2. Predicted and Actual Staff Expenditure, Cessnock 2022-2024 (\$000).


Year	Actual	Predicted
2022	34,034	38,551
2023	35,841	43,260
2024	42,362	47,364

It might be noted that staff expenditures increased significantly in the 2025FY draft financial statements and are now right at the ceiling of what the model predicts. This is unfortunately typical of what we have seen in distressed local governments in the past – there is only so long that a council can run with insufficient staff capacity; eventually matters come to a head and a sudden uplift is almost always the result. Thus, while we understand the recent changes we nevertheless urge renewed vigilance in this area, including potential mitigation through natural attrition where possible.

There are two other major accounting expenditure categories related to efficiency which we also need to examine: ‘materials and contracts’, and ‘other’ (readers might note that ‘depreciation’ refers to past spending on long-lived assets and is thus not relevant to a study of efficiency).

In Figure 6 we chart spending on materials and contracts at Cessnock relative to the peer group for the last four full financial years. It is notable that Cessnock recorded the lowest spending on materials and contracts for the peer group on three of the four years under analysis. This is further evidence of admirable cost control by Councillors and senior staff over a lengthy period of time.

Figure 6. Material Expense per Assessment (\$000)



The last accounting heading that we will look at is 'other expense' which is mainly made up of levies charged by the NSW state government as per Figure 7 which is a cut-out of the most recent audited financial statements:

Figure 7. Other Expense Items.

B3-4 Other expenses


\$ '000	Notes	2024	2023
			Restated
Impairment of receivables			
Other		(465)	(85)
Total impairment of receivables	C1-4	(465)	(85)
Other			
Contributions/levies to other levels of government			
– Emergency services levy (includes FRNSW, SES, and RFS levies)		292	164
– NSW fire brigade levy		500	424
– NSW rural fire service levy		804	804
– Waste levy		4,536	4,361
Donations, contributions and assistance to other organisations (Section 356)		126	(109)
Adjustment for remediation provision estimates		(238)	3,535
Total other		6,020	9,179
Total other expenses		5,555	9,094

In Figure 8 we provide relative data for Cessnock and the peer group. It is notable that Cessnock was in the top quartile for each and every year. Moreover, as Figure 7 makes clear most of these costs are under the control of the NSW state government (or their agencies), not Cessnock City Council. Indeed, some – such as the ESL and fire levies have been increasing at astounding rates of 78 percent or more between 2023 and 2024; well above any measure of inflation. Clearly, the NSW State government have it within their control to alleviate some of the burden on ratepayers – or at least not add to it exponentially – should they wish to do so. {Indeed, if the NSW Government would desist from calculating the waste levy as if Cessnock was a metropolitan council, then some pressure on ratepayers could be mitigated}.

Somewhat confounding the above matters are the provision expenses for remediation works. Typically, most councils tend to under-estimate these expenses and we have noted of late large adjustments as auditors presumably turn their attention to the matter. Provisions are something that are usually estimated by experts in this area – which might include environmental engineers, lawyers, and auditors. There is an extensive note in the draft statements that explains recent movements.

{We note that this result for 'other' expenditure will likely have some implications for the more sophisticated envelopment analysis that follows}.

Figure 8. Other Expense per Assessment (\$000)



This concludes our examination of simple ratio data, that on-the-whole paints a glowing picture of Cessnock City Council's relative spending and cost control over the last four years. In the sections that follow we will instead turn to more sophisticated empirical evidence that provides a nuanced version on matters. It might also be noted that the envelopment analyses that follow are based on a much-expanded cohort of all seventy-one urban councils in NSW⁵, rather than merely the peer group detailed in the first table.

⁵ This is the category that the federal government assigns Cessnock to.

3. Efficiency, 2017-2024

In this section of the report, we will start to employ some of the more sophisticated envelopment analysis. Therefore, it seems to be an opportune time to outline the empirical processes.

Envelopment analysis is a family of linear programming⁶ techniques that allows for the analysis of the efficiency with which multiple inputs are converted into multiple outputs. As such, envelopment analysis is far more consistent with the economic definition of technical efficiency than are the more common single input output ratios we looked at in the previous section. For example, both staff and operational expenditure can be considered as separate inputs in an envelopment exercise, and this allows us to better reflect the various outcomes that are possible through different combinations of production factors. In similar vein, envelopment analysis allows scholars to separate out various proxies for output that better reflect the diversity of goods and services that a local government produces. The specification for the work that we produced makes the advantages of the empirical technique plain:

Staff (\$) + operating expenditure (\$) → residential (no.) + farm (no.) + business (no.) + sealed roads (km) + unsealed roads (km).

Here we consider staff in pecuniary terms to reflect the different skills and productivity that ought to be reflected in remuneration, consistent with Drew, Kortt and Dollery(2015). The output proxies we employ recognise that the respective categories of taxpayers usually have access to vastly different baskets of goods and services. Moreover, we also include as outputs sealed and unsealed roads respectively which properly reflects that these represent the largest items of expenditure, with quite different maintenance schedules (depending on surface). The proxies are thus the best suite to recognise what councils actually do within the limitations of Nunamaker's rule⁷ – and far more realistic than the single outputs used for the earlier ratio analyses. Notably, in the data envelopment analysis (DEA) and free disposability hull (FDH) work that we present, pecuniary data was adjusted to properly reflect the time value of money.

For the work that follows we used an input-orientation consistent with the relevant scholarly literature (Drew, Kortt and Dollery, 2015b). An input-orientation recognises that local government decision-makers have relatively little control over the output proxies, but much more discretion about the resources that they invest into producing same. Otherwise stated, the length of roads is more-or-less given, but how we assign money and staff to maintain them, is certainly something that might change.

⁶ Linear programming is a mathematical technique that can be employed when multiple feasible solutions exist in a mapped function responsive to introduced mathematic constraints. It is iterative in nature and therefore requires significant levels of computing power.


⁷ Nunamaker's rule is a decision-making tool which prescribes that the sum of inputs and outputs ought not exceed a third of the number of decision making units (DMU; that is, local governments). For our seventy-one member cohort our specification is well within this range.

We also used variable return to scale specifications for most of the linear programming (the exception being for the second stage regression work that we do towards the end of this report). This means that we adjusted for the effects of scale.

Envelopment results are both relative and unconditional. Relative means that interpretation of the results can only validly be made with reference to the particular decision-making units and years analysed. Unconditional means that we haven't adjusted for any operating environment⁸ effects (other than size, captured by VRS). We address the unconditional nature of the analyses in our second-stage regression work later in this report.

The most efficient way to understand DEA is generally through a graphical illustration. In Figure 9 we present a simplified input-orientated example. Here the curve drawn between Councils D, B and C represents the theoretically possible efficient frontier. These are the councils that have the best conversion of inputs into a given set of outputs. Councils of this kind are considered perfectly efficient in a relative sense and assigned a score censored at one. Councils in the interior of the curve (such as 'A') represent relatively less efficient decision-making entities. The ratio of the radial distances marked provides a score between zero (perfectly inefficient) and one (perfectly efficient). This number represents the relative technical efficiency of A with respect to the rest of the cohort under analysis (sometimes people multiply this number by one-hundred and then talk about the percent relative technical efficiency).

Figure 9. Input-Orientated DEA



Readers interested in obtaining further information on data envelopment analysis are referred to the seminal works of Cooper et al. (2007) and Coelli et al. (2005).


⁸ Environmental effects and environmental constraints do not refer to climate and the like as is often colloquially the case. Instead, environment here refers to all of the parameters that have an effect on operating costs: demographics, soil substrate, infrastructure burdens, density etc.

It might also be noted that there is some potential for clustering of results, especially if councils face the same harsh decision-making constraints.

In the analysis that follows we use long (seven or eight year) panels for all of the urban councils in NSW. Because of the relatively long timeframe involved we were obliged to use local intertemporal analysis (also sometimes called windows analysis). Local intertemporal DEA is a special kind of moving average which allows us to compare results over time because of overlapping periods. To further assure our results we also bootstrapped⁹ calculations using 10,000 iterations.

Another, slightly different sophisticated approach to measuring efficiency, which we will use in the following section, is called free disposability hull analysis. To ensure that the community receives the most comprehensive picture of relative technical efficiency we also conducted this analysis using the earlier specifications. The main difference between DEA and FDH is that the latter uses a step-wise frontier comprised of the actual results attained by decision making units rather than the curvilinear theoretical efficient frontier. Otherwise stated, DEA tends to be more pessimistic because it compares a given council to an ideal that might not even have been achieved by any of the peer group whereas FDH only compares to what others have actually achieved. Figure 10 provides a graphical comparison of the two approaches.

Figure 10. DEA and FDH Frontier Comparisons.



⁹ Bootstrapping is a probabilistic random re-sampling protocol that is mainly used to reduce potential statistical bias when dealing with a sample.

The mathematic specification for our DEA is:

$$\begin{aligned}
 & \min \theta \\
 & \text{s.t.} \\
 & \sum_{j=1}^n \lambda_j x_{ij} \leq \theta x_{i0}, i = 1, \dots, m \\
 & \sum_{j=1}^n \lambda_j y_{rj} \geq y_{r0}, r = 1, \dots, s \\
 & \sum_{j=1}^n \lambda_j = 1 \text{ (VRS)} \\
 & \lambda_j \geq 0
 \end{aligned}$$

The FDH specification was:

$$\begin{aligned}
 & \min \theta \\
 & \text{s.t.} \\
 & \sum_{j=1}^n \lambda_j x_{ij} \leq \theta x_{i0}, i = 1, \dots, m \\
 & \sum_{j=1}^n \lambda_j y_{rj} \geq y_{r0}, r = 1, \dots, s \\
 & \sum_{j=1}^n \lambda_j = 1 \\
 & \lambda_j \in \{0,1\}
 \end{aligned}$$

Figure 11 presents the local intertemporal data envelopment analysis results for Cessnock relative to various measures of central tendency, for the remainder of the NSW urban council cohort. According to this more sophisticated analysis of efficiency – against what might be theoretically possible – Cessnock has been consistently above average for most of the period.


We note that this more sophisticated robust efficiency evidence differs somewhat to the single input (expenditure) single output (property assessment) work presented earlier. The main reasons for this difference are: (i) the more nuanced specification of different types of assessments, (ii) the important inclusion of roads as an output (separated by surface type), and (iii) the expansion from comparisons to fourteen peers to a comparison of all urban councils (a five-fold increase in peers). The more nuanced specification works against Cessnock somewhat because many urban councils have small cohorts of farmland (please note that the urban classification is dictated by the Commonwealth government schema). For example, the number of farm assessments at Cessnock (893) was well above the median (674). In addition, the inclusion of road lengths also tends to work against Cessnock because of the longer lengths relative to number of properties – it must be remembered that most urban local governments have far more high-density development. For example, Cessnock (955km of roads for 29,312 assessments) has been compared to Strathfield (87kms of roads for 18,218 assessments) – please be mindful that it is not

the length per se that causes distortion, but rather the length relative to number of properties. Furthermore, the expansion of the peer group exacerbates some of our earlier points, but also reduces any potential bias in the sample. In the penultimate section we will also examine the determinants of efficiency which will cast additional light on the operational constraints faced by Cessnock City Council.

It should also be noted that most of the measures of central tendency and spread for the DEA results are clustered towards the upper end of the distribution. Moreover, there is clearly a significant difference between Cessnock, the quartile 1 line, or the lowest relative performer.

We know that Cessnock is doing quite a bit of work to try to improve their relative technical efficiency further within the apparent constraints that they face. In the appendix to this report, we list the efficiency improvements that the council have come up with and provide our comment on same.

Figure 11. Data Envelopment Analysis Efficiency, Local Intertemporal




One of the difficulties with DEA, that we have already discussed, is that it measures a local government against the theoretical production possibility frontier, which may be different from actual results achieved. Accordingly, FDH analysis is also important to provide a more realistic picture of what is going on.

In Figure 12 we plot the more pragmatic FDH for Cessnock against the same measures of central tendency and spread. Under this specification Cessnock performs a little better and notably has had an upwards trajectory in efficiency since 2021. It should also be noted that for the last two years Cessnock has a relative

score of 0.9 (out of a possible 1), despite the unconditional nature of envelopment analysis.

We are optimistic that the commitments made by council detailed in the Appendix might move the relative efficiency up above the median in the future.

Figure 12. Relative Technical Efficiency FDH, Local Intertemporal



4. Tax efficiency, Global Intertemporal.

Underlying much of the regulatory and community dialogue regarding special rate variations is a desire to get maximum 'bang for the buck' (value for the tax dollar). This is an entirely reasonable concern.


We can precisely measure the 'bang for the buck' by conducting FDH and replacing the factors of production with the single input of tax (rates) revenue. Thus, the specification would be:

Total taxation take (\$'000) → residential (no.) + farm (no.) + business (no.) + sealed roads (km) + unsealed roads (km).

For a tax efficiency analysis, it is also important to again conduct variable returns to scale, and bootstrap for maximum assurance (we used 10,000 repetitions).

Figure 13 illustrates tax efficiency over time compared to various measures of central tendency and spread. In most of the years under analysis, Cessnock is below the second quartile. This is a *prima facie* perplexing result, because our Capacity to Pay (CTP) Report clearly showed that Cessnock's total rate take was well-below typical levels. To understand this result people must remember a few facts: (i) that DEA is unconditional, (ii) that comparisons are being made against all urban councils according to the Australian Classification of Local Government schema, and (iii) that a tax efficiency analysis is focussed on the conversion of a single input (tax) into outputs. Thus, while taxes at Cessnock are indisputably on the low side, the conversion of these into outputs may look relatively poor if the environmental constraints at Cessnock mean that outputs consume more tax dollars because of their nature. Our earlier comments regarding road lengths and the like are clearly pertinent here.

Figure 13. Tax Efficiency, 2018-2024 (global intertemporal).



To make matters even more clear, we might consider population density. Thus, Cessnock with a population density of just over 34 people per squared kilometre, is at a significant cost conversion disadvantage when compared to the median (206.6) or third quartile (2,986.13). Economies of density are well-established in the scholarly literature, so these stark differences are clearly going to prove problematic in a naïve unconditional tax efficiency comparison.

Furthermore, population growth has been strongly linked to a deleterious effect on efficiency in the scholarly literature – not only because revenues from growth lag expenditure, but perhaps more importantly because new entrants come with new tastes for local government goods and services. In this regard, it is notable that Cessnock tends to attract people from the Sydney area which on-the-whole has a far wider scope and far higher standard of local goods and services. In addition, new entrants increase heterogeneity which has been shown to reduce efficiency (see, for example, Drew et al., 2024). Notably, population growth for Cessnock is only marginally lower than the third quartile, and almost three times higher than the mean.

In sum, it appears that Cessnock has relatively poor tax efficiency because the particular outputs of Cessnock are so starkly disadvantageous with respect to other councils classified as urban according to the national schema. This has clearly been exacerbated by growth and the heterogeneity that it elicits. Otherwise stated, Cessnock has a higher level of environmental constraint which means that all other things being equal, it will need a higher level of taxation relative to other local government areas in the same classification.

5. The Determinants of Efficiency

In this section we will attempt to empirically identify some of the key determinants of efficiency. This is an important matter to investigate because of some of the results of sophisticated evidence that might have run contrary to *prima facie* expectations; and also in view of the classification of Cessnock by federal authorities as urban, which some people may find odd. To identify the determinants of efficiency for NSW urban councils we conducted second-stage regression analysis – a sophisticated mathematical technique capable of identifying the mean response of a dependent variable (the regressand), to a number of independent variables (the regressors). The regressand for this particular exercise was the constant returns to scale efficiency scores derived from data envelopment analysis according to the following specification:

Staff (\$) + operating expenditure (\$) → residential (no.) + farm (no.) + business (no.) + sealed roads (km) + unsealed roads (km).

Constant returns to scale (CRS) scores were used as the regressand – readers may recall that VRS already controls for size effects and it is clear that using scores of this kind would not have allowed us to test size-related regressors. Against the regressand we tested likely potential determinants as derived from the scholarly literature (see for example, Drew et al., 2015a).

OLS regression was used, with the addition of year dummies to control for the periods under analysis. A fixed effects regression was not suitable given time-invariant (and almost time-invariant) regressors, and a random effects estimate was ruled out by an unfavourable Hausman test. We also included a dummy variable in response to the substantial evidence that amalgamation increased unit costs, *ceteris paribus* (see, for example, McQuestin et al., 2020; Drew et al., 2021; Drew et al., 2023).

The econometric analysis that follows can be specified as:

$$\mathbf{T} = \alpha + \beta_1 \mathbf{P} + \beta_2 \mathbf{X} + \mu.$$

In this specification \mathbf{T} (the dependent variable) is the constant returns to scale technical efficiency score for each council in each year, \mathbf{P} is a vector of relevant population data and \mathbf{X} is a vector of socio-demographic and local government characteristics. μ is an independent identically distributed random error term. It should be noted that natural log transformations were executed where required to correct for skewed distributions, as detailed in Table 2. All standard econometric tests were conducted, and the residuals were confirmed to be near-normal in distribution (a critical assumption for valid statistical reasoning).

Table 2. Definitions of Variables, 2018-2024

Variable	Definition
Rates	
CRS TE	Relative technical efficiency, constant returns to scale
Population	
Pop (ln)	Natural log of the population for each local government area
Pop ² (ln)	The square of the logged population
Density (ln)	Natural log of population density data for each local government area
Controls	
Median employee income	Median employee income (lagged), divided by 1,000
Aged	Proportion of people on an aged pension
Under 15	Proportion of people under the age of 15
DSP	Proportion of people on a Disability Support pension
Newstart (ln)	Proportion of people on a Newstart allowance, logged
Single (ln)	Proportion of people on a Single Parent pension, logged
IPPE (ln)	Natural log of the carrying value of infrastructure (\$'000)
Year	A dummy variable to control for the effect of different years
Amalgamation	A dummy variable to control for whether or not a council was amalgamated in 2016

We will not table the coefficients, standard errors and statistical significance of each regressor because recent experience has demonstrated to us that most end users find this very confusing. Instead, we will only discuss the sign of important variables which have met conventional statistical significance thresholds. We note that for the model overall it met all relevant statistical tests and had an appropriate level of explanatory power.

Only two relevant variables were statistically significant overall for urban councils, and both of these were at the highest level of certainty (the 1% level): aged persons and the proportion of people identifying as Aboriginal or Torres Strait Islander (ATSI).

For the urban cohort over the seven-year panel we found strong evidence that as the proportion of aged people increased, CRS efficiency decreased. This is problematic

for Cessnock because its aged cohort¹⁰ numbered some 11.22 percent over the period, whilst the median for the remainder of the peer group was 9.64 percent, and the mean was 10.45 percent. Moreover, councils like Cessnock are disadvantaged twice – not only does this variable reduce the potential for efficiencies, but it also reduces revenue significantly as detailed in the *Capacity to Pay Report*. Otherwise stated, having a higher than typical aged pension cohort is clearly going to have an important deleterious impact on financial sustainability.

A higher proportion of people identifying as ATSI was also found to have a significant deleterious effect on efficiency for the entire cohort, consistent with the scholarly literature. This is important because the proportion of people at Cessnock who identify as ATSI was 10.2 percent according to the most recent census, compared to the median for the local government cohort of 3.65 percent, and a mean of 4.08 per cent. Indeed, the result for Cessnock was far higher than even the third quartile at 6.5 per cent.

Before closing this section, we would like to clarify that it is typical to have only a few statistically significant variables for a second-stage panel regression of FDH efficiency. This does not mean that other variables are not important, simply that we cannot say so, with respect to the whole urban cohort, with statistical certainty. However, the two variables that we can identify as unambiguous determinants of urban local government efficiency certainly suggest higher environmental constraints at Cessnock relative to the peer group. This underscores our earlier comments regarding the achievements of management for the local government area.

6. Recommendations

IPART and the community have a reasonable expectation that as part of the SV process steps will be taken to make council as efficient as possible. We have previously shown that cost control is exceptional at Cessnock, and also that environmental constraints are formidable. However, there is always more that can be done – notwithstanding that our foregoing comments mean that future efficiency enhancements are likely to be immaterial to an SV.

Council staff have been working on a list of efficiencies that we include in the appendix. Part of the important work that we do is to assure IPART and the community around this and other matters (such as the long-term financial plan (LTFP)). Professor Drew has examined the list carefully with respect to both plausibility and materiality. He has also contacted managers to ask questions and provide guidance around the list. Furthermore, he has worked with senior management to explore a range of potential efficiencies drawing on both scholarly knowledge and his experience working with dozens of other councils. Otherwise stated, the list in the appendix has been assured and is thus a reliable indication of what council will be able to achieve.

¹⁰ People receiving an aged pension.

In addition to this critical (albeit largely unseen) task described above we also provide our own list of efficiencies that we would encourage Councillors and management to implement as quickly as possible.

1. Future spending: reconsider future spending plans to remove some 'discretionary' items (the LTFP has already been adjusted to reflect this). Ensure all future spending decisions are made with specific reference to the LTFP, and only proceed if already included, fully funded, or matched with commensurate expenditure cuts. Carefully consider the wisdom of proceeding with discretionary spending even if putatively fully funded through grants – for example, a lot of the grants funding since COVID has been for non-core infrastructure that has entrenched fiscal illusion¹¹ and inflated unit costs going forward and might not have been prudent in hindsight.
2. Ensure that any future discretionary spending is supported by random surveys citing whole-of-life cost.
3. Reprice non-regulated fees and charges as quickly as possible. This is not just about additional revenue – which is clearly needed – but also orientated to improving equity and reducing consumption to economically efficient levels.
4. Related to the above, make greater efforts to more promptly recover outstanding rates, fees and charges. In our capacity to pay report we showed that councils with much higher revenue efforts had far lower outstanding rates and charges. It is essential that people pay for what they consume if we want economically efficient levels of consumption and effective price signals. Notably, IPART also requires us to report on our strategy related to this matter.
5. Carefully control staff expenditure with reference to the modelling and commentary provided in this report.
6. Actively encourage staff to use their outstanding level entitlements – especially long-service leave.
7. Introduce a robust tracked staff suggestion scheme – likely online – that escalates to decision makers. Non-cash proportionate incentives and a reasonable expectation that good ideas will be acted on, or where not possible that reasons for eschewing will be provided. We acknowledge that a scheme is in place, but not as robust as we envision.
8. Related to the above introduce a similar tracked and traced staff suggestion system for reporting potential WHS risks proactively. We acknowledge that a scheme is in place, but not as robust as we envision.
9. Make greater efforts on community education. The SV process is as start as are the videos we have been producing (and other videos that can be accessed from Professor Drew's YouTube site). However, we also need to be conveying more factual information to community in ways that they are likely

¹¹ Fiscal illusion is the economic term to describe the situation whereby citizens struggle to understand the true cost of the things that they consume, and by extension, the true state of financial sustainability. For example, if people see a new playground built they usually don't recognise that it might have been fully funded by a grant – instead they believe that their council must be doing well (are sustainable) and have more than sufficient revenues.

to consume it. For instance, rates notices, should include facts on revenue, spending and the like and these should change each quarter. A second related example is price signalling – all receipts should include the cost paid by the resident as well as the quantum of the subsidy provided by council, where relevant. This includes things like swimming pool entry costs, library book borrowing receipts and the like. We simply cannot expect residents to understand the financial sustainability situation of council if we do not provide them with information.

On financial sustainability more generally, Prof Drew has offered to work with Councillors for a day long (*pro bono*) workshop in the near future. In addition, the following recommendations seem apposite:

10. Councillors and senior (non-finance) staff should consider enrolling in an appropriate financial sustainability training course. UoN used to do day long courses of this kind and the UNE may offer similar in the future. Other providers might also exist.
11. Encourage neighbouring councils to also do the above course – there is a strong policy risk of future structural reforms. It would be devastating for the community at Cessnock to work hard on their own sustainability journey, only to be amalgamated with a neighbouring community who hadn't even taken the first steps.

We remind end-users of our assurance work for council originated efficiencies in the appendix. We commend this report to IPART and the community.

References

Australian Bureau of Statistics (ABS) (2022). *Regional population methodology*. Available at: <https://www.abs.gov.au/methodologies/regional-population-methodology/2021-22>

Bevan, G. and Hood, C. (2006). What's Measured is What matters: Targets and Gaming in the English Public Health Care System. *Public Administration*, 84(3): 517-538.

Bozeman, B. (2019). Public Values: Citizens' Perspective. *Public Management Review*, 21(6): 817-838.

Coelli, T., Rao, D., O'Donnell, C. and Battese, G. (2005). *An Introduction to Efficiency and Productivity Analysis*. Springer: Singapore.

Cooper, W., Seiford, L. and Tone, K. (2007). *Data Envelopment Analysis*. Springer: Singapore.

Dollery, B. E., Wallis, J. L. and Allan, P. (2006). The Debate That Had to Happen But Never Did: The Changing Role of Australian Local Government. *Australian Journal of Political Science*, 41(4): 553-567.

Drew, J. (2021). *Saving Local Government*. Springer Palgrave: Singapore.

Drew, J. (2025). *Creating Human Value as Public Management Theory*. Routledge: London

Drew, J. and Dollery B. (2014). Keeping It In-House – Households as an Alternative Proxy for Local Government Output. *Australian Journal of Public Administration*, 73(2): 235-246.

Drew, J. and Dollery, B. (2015). Less Haste More Speed: The Fit for Future Reform Program in New South Wales Local Government. *Australian Journal of Public Administration*, 75(1): 78-88.

Drew, J. and Dollery, B. (2017). The Price of Democracy? Political Representation Structure and Per Capita Expenditure in Victorian Local Government. *Urban Affairs Review*, 53(3): 522-538.

Drew, J., Kortt, M. and Dollery, B. (2015a). Peas in a Pod: Are Efficient Municipalities Also Financially Sustainable? *Australian Accounting Review*, 77(26): 122-131.

Drew, J., Kortt, M. and Dollery, B. (2015b). No Aladdin's Cave in New South Wales? Local Government Amalgamation, Scale Economies and Data Envelopment Specification. *Administration & Society*, DOI: 10.1177/0095399715581045.

Drew, J., McQuestin, D., and Dollery, B. (2021). Did Amalgamation Make Local Government More Fit for the Future? *Australian Journal of Public Administration*, DOI: 10.1111/1467-8500.12530.

Drew, J. McQuestin, D, and Dollery, B. (2023). Fiscal Outcomes Arising from Amalgamation: More Complex than Merely Economies of Scale. *Public Management Review*, <https://doi.org/10.1080/14719037.2023.2174586>

Drew, J., Miyazaki, M., and McQuestin, D. (2024). Does Inequality Still Matter? Income Heterogeneity and Local Government Expenditure. *The Economic Record*, Accepted 17/4/24. doi: 10.1111/1475-4932.12812

Drew, J., Razin, E. and Andrews, R. (2018). Rhetoric in Municipal Amalgamations: A Comparative Analysis. *Local Government Studies*, doi:10.1080/03003930.2018.1530657.

Fenwick, C. (1920). Democracy and Efficient Government – Lessons of the War. *American Political Science Review*, 14(4): 565-586.

Ferguson, C. E. (1972). *Microeconomic Theory*. Richard D. Irwin: Homewood Illinois (Third Edition).

Friedman, M. (1993). *Why Government is the Problem*. Hoover Essays: Stanford.

McQuestin, D., Miyazaki, M. and Drew, J. (2020). Do Amalgamations Make a Difference? What We Can Learn from Evaluating the Policy Success of a Large Scale Forced Amalgamation, *Public Administration Quarterly*, 45(3): 278-298.

O'Flynn, J. (2007). From New Public Management to Public Value: Paradigmatic Change and Managerial Implications. *Australian Journal of Public Administration*, 66(3): 353-366.

Appendix 1 – Efficiency Assurance

Early in October Professor Drew spent a number of days evaluating the list of proposed efficiencies produced by council. Lengthy conversations were had with responsible managers and the proposals assessed according to the criteria of: plausibility, meaningfulness, quantifiability and deliverability. Some figures were altered significantly after investigation. To be more specific, we reduced the value of some efficiencies and removed other potential savings because sufficient evidence or assurance could not be had. Thus, what has been assured, in the Table below, will probably transpire to be a light under-estimate of what might be delivered. All assured efficiencies have been imputed into the LTFP and management understand that they must be delivered and that they will be held accountable for doing so.

Indeed, we expect actual efficiencies to improve even further in response to councillor backing for a tracked efficiency suggestion scheme that will be escalated to decision-makers, with feedback given to the staff proposing the potential efficiency, and a non-cash incentive provided. We plan to leverage off existing technology – a little used programme that can be adapted to our efficiency needs. Key to this scheme will be automatic escalation until a decision and feedback are provided.

The above comments on efficiency must be understood in terms of this report, especially the FDH analysis and analysis of determinants. Cessnock was already a relatively efficient operation as per the robust evidence, so in this context the scale of assured efficiencies probably surprises on the upside.

Initial efficiency estimates of some \$1.5 million were included in the Fact Sheet, but this number improved as more work was conducted in the area. Efficiencies were also discussed in some detail during the public information sessions.

Notably the aforementioned efficiencies excluded the five-year freeze on real staff expenditure which is crucial to the successful delivery of the LTFP. Further details are available in council documents.

Summary Table of Assured Efficiencies (more details in internal documents)

Item	Saving	Revenue Increase	Progress
Finance Efficiencies	105,000	10,000	
Records and Council Services	9,100	17,775	
Regulatory Services		194,000	
Community Services	319,793	18,000	
Communications	7,295		
People & Culture	22,249		

Internal audit	50,710		
Development Services	0	0	
Airport		397,000	
Waste	140,000	117,000	
IT	40,000		
Customer Service	5,000		
Infrastructure	66,000		
Parks and Maintenance	756,453	135,000	
TOTALS	1,521,600	888,775	
GRAND total (ADDED)			2,410,375