

**ALCOHOL, TRAUMA CASES, CAUSAL INFERENCE AND POLICYMAKING:
A BRIEF REVIEW OF THE ANALYSIS FOR DISTELL BY MURRAY AND MCGORIAN¹**

Dr Seán Mfundza Muller

26 April 2021

Requested by: the DG Murray Trust²

Summary

The statistical analysis used in the Distell study is very low quality and, for reasons relating to quality, method and data availability, insufficient to justify key assertions made in the study itself and by others with reference to the study. The study reflects a primitive and inadequate understanding of causal inference and how evidence informs policy. Claims that the study has credibly shown lockdown restrictions, or effects of lockdowns on mobility, to be a greater causal factor in trauma admissions than alcohol bans/restrictions are false given the study's numerous weaknesses.

The study is correct in stating that the graph (of trauma admissions to Western Cape public hospitals) the authors reference is insufficient to infer a causal relationship between alcohol bans/restrictions and trauma cases. That does not in-and-of-itself demonstrate that such restrictions cannot be justified, but rather that more evidence (than the graph) is required to support the use of alcohol bans. Contrary to the study's claim, however, such evidence need not only be quantitative but could be a mix of the best available quantitative and qualitative information and analysis.

I was not provided with any such additional data or information and therefore cannot take a position on the actual question at hand, namely whether alcohol bans/restrictions do reduce trauma admissions to public hospitals. I would encourage the government to release the data it has and the rationale for its position. The authors of the Distell study should also publish the data they used and all details of their statistical analysis.

¹ The review here refers to the document, "A deep dive into the relationship between trauma admissions and lockdown measures during the COVID-19 pandemic in South Africa. A HIGH LEVEL OVERVIEW" which is six pages long. Neither the authors nor Distell appear to have made any further information available, including the data used or the detailed statistical results.

² Given, legitimate, concerns about the credibility of research commissioned by parties with a vested interest (such as the research produced by Murray and McGorian on behalf of an alcohol industry body), the author conducted this analysis *pro bono* and, furthermore, did not have any prior view on the merits of either alcohol bans or claims as to their effects.

Among his qualifications to conduct this review are: a PhD on the use of randomised trials in economics to inform policy decisions; publications on the limitations of randomised experiments, the challenge of external validity (extrapolating from causal estimates to other contexts), and the link between correlations and causal relationships; extensive training in microeconometrics at the University of Cape Town and University of Oxford; publications and ongoing work on methodological and philosophical questions pertaining to the use of evidence for policy decisions, including two recent papers on Covid-19. One of these Covid-19 papers criticises national government decisions regarding the timing and use of lockdowns, as well as the absence of a coherent, well-justified long-term Covid-19 strategy for the country. The other argues that a strong case can be made for mask usage as a precautionary preventive measure even in the absence of adequate evidence from randomised evaluations.

The study by Murray and McGorian for Distell (henceforth 'the Distell study' or 'the study') makes some correct assertions, but is riddled with weaknesses and errors that render it an unreliable basis for claims about the relationship between alcohol bans/restrictions and trauma admissions to public hospitals. It is also an unreliable basis for claims about the associated policy issue. The study contains a number of the errors similar to those it seeks to criticise and many more in addition to those.

1. What the Distell study gets right

The Distell study argues that it is problematic to infer a causal relationship simply from an association; this is correct. In some places, it argues that therefore more is needed to justify the government's stance than the kind of association shown in the graph of trauma admissions to Western Cape public hospitals the study rests upon; this is also correct. The study suggests that evidence of reductions in trauma admissions in other countries without alcohol bans is suggestive that (in my phrasing) the absolute value of the apparent negative association between alcohol bans/restrictions and trauma admissions may be upward biased; this is a legitimate hypothesis.

The study argues that such reasoning is the basis of the government's view that alcohol bans reduce trauma cases, from which it would follow that the government's position is based on inadequate evidence and reasoning. However, the study does not demonstrate that this is, in fact, the only basis for the government's position. A great deal more information and argument is required to justify the claim made.

In that regard, the government (national and provincial) ought to release its data and reasoning in the interests of transparency.

2. What the study gets wrong

If the study had limited itself to arguing that the evidence in the public domain is insufficient to support the government position, it would have been unobjectionable – albeit also not particularly interesting since these points are quite obvious. However, the Distell study seeks to go further and conducts what it claims to be a credible statistical analysis. Upon that basis it goes further and make new claims about the causes of trauma admissions. Finally, the study commits itself to a particular position on the nature of evidence, and the link between evidence and policy. All of these aspects of the study are deeply flawed.

2.1. Flawed claims about evidence and policy

The Distell study relies on a very primitive understanding of causal inference. First, it implies a causal relationship of the kind in question here can be established 'with certainty'; that is false in almost all cases of population-level social science research. At best, we can establish such relationships exist with a high degree of confidence. Second, the study states that establishing causality can only be done through randomised evaluations. While there are proponents of that view, it also has many critics and in recent years its proponents have backed-away from strong versions of the associated epistemic hierarchy. Third, while the study correctly characterises the basic conceptual problems of spurious correlations and omitted variables, it fails to appreciate that actual empirical analysis of the kind it attempts cannot be framed in such simple terms. (More on this in the discussion of statistical analysis below). Fourth, the notion of counterfactual expressed is too crude; the study fails to make reference to approaches such as interrupted time series analysis and difference-in-differences that rely upon particular assumptions in order to estimate causal effects. Indeed, the authors seem entirely unaware of the extensive econometric literature on causal inference under different sets of identifying assumptions.

Added to this is a correspondingly primitive notion of what constitutes ‘evidence’. Aside from dogmatic proponents of randomised trials, most serious texts on the role of evidence in policymaking recognise that evidence may come in many forms. Descriptive statistics are one form. So too are qualitative observations of, in this case, hospital staff. Indeed, in most policy contexts randomised evaluations are unavailable: so a reliance on a narrow conception of evidence combined with a rigid evidence hierarchy would mean that policymakers would be unable to make decisions in most situations. That is exacerbated in circumstances such as Covid-19 where there is limited opportunity to gather evidence and design studies. Much of the underlying reasoning of the study is flawed in this respect and therefore so too are the associated policy claims.

2.2. Flawed statistical analysis

The authors claim to have conducted a statistical analysis, however they do not report the approach to the analysis or the findings in a credible manner. The specification of the statistical/econometric models is exceedingly crude and poorly described. The characteristics of the data itself are not described. And the statistical results are not properly reported.

Even putting such egregious omission of information aside, it is evident that the statistical analysis is far too primitive. The first flaw is in relation to what correlation between potential explanatory variables implies for appropriate regression specification. It is correct to say that correlation between two possible explanatory variables may reflect omitted variable bias if one is left out of a regression (or other form of) analysis. However, it does not follow that merely including the other variable resolves the issue since there are a variety of ways in which the variables could be related – a point that can be illustrated easily using causal graphs. Any causal relationship between variables may also involve other excluded variables. At best, the study’s approach only calls into question the extremely naïve use of a univariate regression; however, it does not provide any evidence that the government has in fact utilised such a model for its decisions.

However, because of the apparent lack of understanding of the more complicated considerations involved, the study then proceeds to conduct an analysis which purports to be a credible assessment of the causal relationship between lockdown restrictions, alcohol bans/restrictions and trauma admissions. Except that it is nothing of the sort: taking an extremely bad/naïve model and adding one other variable does not make for a much more credible model. Given the use of observational data, the specified model needs to bear sufficient resemblance to the underlying data generating process (DGP) in order to have any confidence that the estimated parameters, reduced form or otherwise, are unbiased. The study provides no discussion of what the DGP might look like and therefore makes no case at all for the credibility of the models that were apparently estimated. This applies both to the comparison of alcohol bans/restrictions and lockdowns, as well as alcohol bans/restrictions and mobility restrictions (as measured by Google mobility data).

A third issue concerns some more specific aspects of the model specification. A first example is that the authors state that they used a “one period autocorrelation adjustment”. Numerous questions arise. Why one period? Were other specifications estimated? Are the results robust to this assumption? And so forth. A second example is the mobility ‘analysis’. The authors limit themselves to regressions containing two variables – using either mobility or lockdowns – but do not explain why they do not consider a specification with all three as explanatory variables. A third example relates to the study’s attempt to examine the specific phenomenon of weekend binge drinking and whether this can shed light on the possible effects of alcohol bans/restrictions. The authors create a new dependent variable to represent differences between week and weekend trauma cases, but this seems like a clumsy approach. For example, an alternative would be to retain trauma admissions as the dependent variable, include a binary variable representing week and weekend days, then include interaction

terms between the binary variable and other explanatory variables. Putting aside the above point that the overall approach to specification is flawed, the study says nothing about the robustness of its results to alternative specifications – contrary to good practice.

The time-series nature of the data, which is given no serious consideration in the study, compounds all the above weaknesses.

For these and other reasons it is no exaggeration to say that the statistical analysis lacks any credibility whatsoever. It also fails to consider subtle issues, such as possible behavioural changes over time in response to alcohol bans/restrictions which could be important for a serious analysis of their effects.

2.3. Flawed/unwarranted claims

The final major problem in the study concerns its flawed and/or unwarranted claims.

Given the limited understanding of the role of evidence in policymaking evinced in the paper, the claims about what would constitute adequate evidence are poorly-founded. And in the absence of more information about the basis for the government's position, the authors cannot assert that the position is poorly-founded merely on the basis of one graph that they have chosen to rely on. Indeed, the authors make no attempt to even demonstrate that the graph is representative of the evidence base for the government position – contradicting basic principles of good research.

The limited data available and associated limitations of any statistical analysis mean that any responsible analysis could say very little about the question at hand. The many weaknesses and flaws in the authors' approach compounds this. Yet the study makes a number of bold claims. One is that lockdowns and mobility restrictions had a greater effect on trauma admissions than alcohol bans/restrictions. Another is that once such factors are taken into account, alcohol bans/restrictions have no effect at all. A third is that an alcohol ban has limited benefit in addition to a curfew. These claims are effectively baseless, because of the flaws in the statistical analysis.

The study claims to 'refute' the government's position; this is false. At best, it raises concerns about the government's apparent position on the basis of a lack of public information and the simple issues mentioned in 1. above.

3. Conclusion

The Distell study does *not* demonstrate that alcohol bans/restrictions do not play a role in reducing trauma admissions. It also does *not*, credibly, demonstrate that lockdowns play a greater role than such measures. This is because the methods employed are hopelessly crude, and indeed suffer from more sophisticated versions of the very same criticisms the study deploys against the claimed justification for alcohol bans/restrictions.

Whether alcohol bans/restrictions reduce trauma admissions in South Africa, and to what degree, remains an open, empirical question. As a conceptual matter, it is also important to note that the effect could vary across context and across time, which further complicates efforts to estimate it. Nevertheless, governments often do not have detailed or definitive statistical/econometric evidence for decisions: what matters is the preponderance of evidence and knowledge, of various kinds, which ought to be set out accordingly along with the likely costs and benefits of different decisions.

Government should make the evidence and rationale for its position public: thereby ensuring a properly-informed public debate and limiting the scope for flawed, misleading analysis of the kind presented in the Distell study.