New modelling of alcohol pricing policies, alcohol consumption and harm in Scotland

Summary of key model developments

This briefing note presents an overview of the most significant differences between the model used for a new report from the Sheffield Alcohol Research Group on alcohol policy in Scotland, commissioned by Scottish Government, and the model used in our 2016 report on Minimum Unit Pricing.

In 2016, the Sheffield Alcohol Research Group published analysis of the potential impact of introducing a Minimum Unit Pricing (MUP) policy for alcohol in Scotland. This analysis used the Sheffield Alcohol Policy Model (SAPM). We have now published new analysis that appraises:

- The potential impact of changing the MUP threshold in Scotland from its current level of 50p per unit
- The impact that changes in drinking behaviours during the COVID-19 pandemic could have on long-term alcohol consumption and harm
- The impact of the alcohol duty reforms introduced in August 2023
- The impact of recent high inflation on the real-terms value of the MUP threshold and how alternative approaches to linking the MUP to inflation in the future might address this.

These new analysis use a new model, the Sheffield Tobacco and Alcohol Policy Model (STAPM). Further details of STAPM can be found in the new report, as well as on the STAPM webpage and in technical reports on the modelling of alcohol consumption and harm and tax and economic outcomes.

STAPM uses many of the same data inputs, evidence and assumptions as the previous, SAPM, model and shares many aspects of its structure. However, there are some key differences between the two models which affect how their results should be interpreted. These are addressed in detail in the resources linked above, but we summarise the two most significant here:

1. **Drinker group definitions**

In SAPM, all modelled individuals were assigned to a drinker group (moderate, hazardous or harmful) based on their alcohol consumption *at baseline* (i.e. prior to any policies being introduced). This meant that if their drinking levels changed as the result of a policy, such as MUP, being introduced, their assigned drinker group would not change. For example, somebody drinking 16 units per week at baseline would be categorised as a hazardous drinker. If a policy was introduced that caused their consumption to reduce to 13 units a week, meaning they are now drinking at moderate levels, they would still retain their initial classification as a hazardous drinker for all reporting in the model.

In STAPM, each individual's drinker group classification is updated as their alcohol consumption changes. This means that the same drinker would be reclassified as a moderate drinker from the point that their drinking fell to 13 units onwards. The reason for this updating of drinker group classifications over time is because the STAPM model takes a dynamic

approach to modelling alcohol consumption at baseline, based on a simulation of changes to individual alcohol consumption over the course of their lives.

This change in approach means that it is now possible to report changes over time in the number of people drinking at moderate, hazardous or harmful levels. For example, we can now estimate how raising the MUP level to 60p per unit would impact on the number of people drinking at harmful levels. However, it means that reporting outcomes for the average amount of alcohol consumed by drinker group from STAPM gives results that may initially seem puzzling compared to the 2016 report. This is because the comparator against which change is assessed is different between the two models. In SAPM, change in average alcohol consumption is assessed by comparing the alcohol consumption of the same set of individuals between a fixed baseline year and a certain year after the policy change. This comparison is made separately for each drinker group. In STAPM, change in average alcohol consumption is assessed by comparing the alcohol consumption of the individuals who are assigned to each drinker group in a particular year of the simulation between a "business as usual" scenario (which is the dynamic baseline used for comparison) and a "policy effect" scenario. The comparison in STAPM is therefore made between different sets of individuals, which means that the reporting of outcomes for average alcohol consumption by drinker group is influenced by changes to the alcohol consumption of individuals who remain within each drinker group and by changes to drinker group membership.

For example, in STAPM, increasing the MUP level to 60p per unit leads to the observation of a larger relative reduction in average alcohol consumption across the whole population (-6.7%) than within any of the three drinker groups individually (-0.3%, -1.1% and -2.6% respectively for moderate, hazardous and harmful drinkers). This can occur because some of the harmful drinkers who reduce their consumption will be reclassified as hazardous drinkers. It is most likely that the harmful drinkers who move into the hazardous drinking group are those drinking at the lower end of the harmful spectrum (i.e. close to the hazardous/harmful boundary). As a result of this change in group membership, the average alcohol consumption in the harmful drinker group will rise (because people drinking at relatively low levels in this group have left) and the average alcohol consumption in the harmful drinker group will also rise (because new individuals drinking at relatively high levels for this group have entered). The effects of changes to group membership will partially offset any fall in average consumption arising from policy effects causing individual drinkers to drink less. A similar effect at the moderate/hazardous boundary leads to this counterintuitive finding that average alcohol consumption at the overall population level is observed to fall to a greater extent than the average alcohol consumption in any of the drinker groups.

For health outcomes, such as mortality, similar processes apply that can lead to the reported changes to the number of deaths for people in a certain drinker group being different to what might be expected. This difference is again due to individuals moving between drinker groups. For example, if a lot of individuals move down from the hazardous to moderate drinker group, then this could increase the number of moderate drinker deaths because the moderate drinker group is now significantly larger. In addition, since the risks associated with alcohol consumption can remain elevated for several years after an individual reduces their alcohol intake, individuals moving to a lower drinker group will carry additional risk with them as a result of their previous drinking history. This risk carried forward from previous drinking can also be a factor that increases the number of deaths in the lower-level drinker group that the individual has moved into.

For these reasons, although we continue to present results by drinker group, we have included these in an appendix to the new report in order to reduce the potential for confusion.

2. All-cause mortality

In both SAPM and STAPM, alcohol consumption is linked to mortality risks of 45 separate alcohol-related health conditions. As it is impossible to prevent death indefinitely, only postpone it, it is important to consider what happens in the future to individuals for whom an alcohol policy reduces their drinking and so prevents them dying due to alcohol. In SAPM, changes in *alcohol-attributable* deaths were reported. Under this approach, a policy which reduces alcohol consumption might avert deaths from an alcohol-related cause such as liver disease. However, some of those same individuals who no longer die from liver disease may subsequently die from other causes within the time horizon of the model. SAPM limits its reporting to whether those individuals will subsequently die from any other disease related to alcohol, rather than reporting the total change to the number of deaths from all causes.

As the aim of effective public health policy is to improve health and extend life, in STAPM we take an all-cause mortality perspective. This means that there is no distinction between delayed deaths from alcohol-related or unrelated causes. We simply measure the overall number of deaths in the population from any cause and compare these between modelled scenarios. As this approach on its own does not capture the benefits of extending life, we present our mortality estimates alongside estimates of the total number of Years of Life Lost (YLLs) to premature death. This YLL measure multiplies the change to the number of deaths by the expected remaining lifespan of the individuals concerned. This measure therefore captures the benefits of extending lifespan within the time horizon of the model over and above the number of deaths.