Double-Counting Lives Saved in the Vaccine Impact Modelling Consortium: Analysis and Proposed Solutions

Executive summary

Combining estimates of deaths averted across pathogen-specific models may lead to double-counting the benefits of Gavi-supported vaccinations. This report describes why and how this problem arises, which is due to persons who could be fatally infected with more than one pathogen of interest. An analogous situation is the population attributable fraction of a specific risk factor for some disease, which can appear to sum to more than 100% of cases averted when combining estimates for multiple risk factors. Drawing from the rich epidemiologic literature on the population attributable fraction, some methods for reducing the problem of double-counting are proposed.

Description of the problem

The Vaccine Impact Modelling Consortium (VIMC) estimates the impact immunizations against mortality from ten pathogens in low- and middle-income countries (LMICs). For this purpose, research teams within the VIMC use mathematical or statistical models of disease transmission/burden and vaccination. Vaccine impact on mortality is estimated by subtracting the predicted disease-specific deaths assuming some number of vaccine doses as model inputs from the predicted deaths under counterfactual vaccine coverage assumptions.

The models used by the VIMC research teams have been developed independently from each other and generally for purposes other than estimating the impact of Gavi-supported vaccination. In practice, this means that VIMC impact estimates for each pathogen are calculated independently from each other pathogen, with the same demographic inputs used for each pathogen-specific model. This use of separate models leads to the concern that the VIMC is double-counting the number of deaths averted by Gavi-supported vaccination across pathogens.

Averted deaths may be double-counted for any persons who could be fatally infected with more than one targeted pathogen during the time periods of interest. Informally, these are persons who, in the absence of vaccination would die from one pathogen, but in the presence of vaccination against that pathogen would later die from a different pathogen.[1] To illustrate, consider vaccination against *Haemophilus influenzae* type b (Hib) and against measles virus. In the absence of any vaccination, a hypothetical child might die of Hib meningitis at 3 months of

[1] More formally, members of the population can be considered to belong to one of a number of "principal strata" describing the infections they would encounter in the presence/absence of different counterfactual vaccination programs: never fatally exposed to any pathogen, fatally exposed to measles only, fatally exposed to Hib only, measles averted and then exposed to Hib, Hib averted and then exposed to measles, and so on for all combinations of the 10 pathogens.

age. In the presence of Hib vaccination but not measles vaccination, that child might be protected against Hib but go on to die of measles at 15 months of age.

In this toy example, comparing either "Hib vaccine only" or "measles vaccine only" against a "no vaccine" counterfactual would not yield a death being averted, since the child would still die from the other pathogen. Comparing "Hib and measles vaccine" against a "no vaccine" counterfactual would yield one death being averted, as the child would die in the absence of vaccination but would survive in the presence of both. However:

a) Comparing "Hib and measles vaccines" to a "measles vaccine only" counterfactual would yield one death averted, as the child would die of Hib without Hib vaccine but would survive in the presence of both;
b) Comparing "Hib and measles vaccines" to a "Hib vaccine only" counterfactual would also yield one death averted, as the child would die of measles without measles vaccine but would survive in the presence of both.

If deaths averted are combined from individual vaccine models that include other vaccines as counterfactuals, then deaths averted will be counted multiple times for some population members. This is the situation within the VIMC, because the individual pathogen-specific models make the implicit assumption that other vaccines are in use. Within the pathogen-specific VIMC simulation models, cause-specific mortality is only tracked for the pathogen of interest; mortality due to other pathogens is lumped into overall age-, time-, and country-specific death rates. These death rates are based on UN population data and projections, which are in turn the result of various actual events in the real world, including vaccination programs as actually implemented. This means that the "no Gavi support" counterfactual for a specific pathogen assumes (implicitly) that all other vaccination programs occur at the actual Gavi-supported level.

Combining deaths averted from pathogen-specific models in the VIMC thus leads to double-counting of deaths averted when considering vaccination as a whole. In the next section, we consider another common situation in epidemiology, combining population attributable fractions (PAFs) which faces a similar double-counting problem. We then conclude by using the analogy to PAFs to suggest methods for reducing double-counting when estimating deaths averted in VIMC.

An analogy to population attributable fractions

The PAF is a measure that quantifies the degree to which a specific risk factor contributes to the overall disease burden in a population.[2] The PAF represents the proportion of disease incidence that could be eliminated if the risk factor were removed. As an epidemiologic measure, the PAF was first used in the literature over 60 years ago, and numerous papers have discussed the

[2] Poole C. "A history of the population attributable fraction and related measures." *Ann Epidem* 2015; 25(3): 147-54

utility, limitations, and interpretation of PAFs since that time. One of its well-known properties is that PAFs for multiple disease risk factors can sum to greater than 1.0, which implies that more than 100% all disease incidence could be eliminated if the risk factors were removed.

PAFs can sum to more than 1.0 because some at-risk individuals may have multiple risk factors sufficient to cause disease. For example, both hypertension and cigarette smoking are independent risk factors for coronary artery disease, with estimated PAFs for the United States population of 0.25 and 0.22, respectively. If cigarette smoking and hypertension could be eliminated, we would not expect to prevent (25% + 22%=) 47% of incident cases of coronary artery disease, because some persons with hypertension also smoke cigarettes and could develop coronary artery disease from either exposure. The proportion of incident cases prevented by removing both cigarette smoking and hypertension will be something less than the sum of the individual PAFs; how much less depends on the degree of overlap between smoking and hypertension. Summing PAFs for individual risk factors is analogous to the problem of double-counting deaths averted within the VIMC. In both situations, the overlap of risk factors in some members of the population leads to over-estimation of the effects of combining interventions against individual risk factors.

Within the literature on PAFs, several methods have been proposed for combining PAF estimates from individual risk factors to estimate the combined effect of removing multiple risk factors. The simplest approach is to take the product of the complement of the individual PAFs:

$$PAF_{combined} = 1 - \prod_{r}(1 - PAF_r)$$

where *r* indexes each risk factor.[3] This method is computationally simple and does not require data on the population of interest beyond the individual PAF estimates for each risk factor. However, it assumes that risk factors are both biologically independent and randomly distributed. It is also an approximation that will approach 1.0 only asymptotically, even if all true risk factors are included.

An alternative to combining PAFs is to consider removal of each risk factor sequentially, rather than assuming they are removed simultaneously. The sequential attributable fraction (SAF)[4] for a risk factor *k* can be calculated as:

$$SAF_k = \sum_{r} \frac{C_r - C_{r(-k)}}{C}$$

where C is the total number of cases, C_r is the cases occurring among persons with risk factor *r*, and $C_{r(-k)}$ is cases occurring among persons with risk factor *r* after *k* has been eliminated. These

[3] Zapata-Diomedi B, Barendregt JJ, Verrman JL. "Population attributable fraction: names, types and issues with incorrect interpretation of relative risks." *Br J Sports Med* 2018; 52(4): 212

[4] Rowe AK, Powell KE, Flanders WD. "Why population attributable fractions can sum to more than one." *Am J Prev Med* 2004; 26(3):243-9

"sequential attributable fractions" (SAFs) do account for overlapping risk factors, in that the SAF for removing an individual risk factor is conditional on which other risk factors have already been removed. The SAF also does not require biological or statistical independence between risk factors. However, calculating the SAF for a specific risk factor is "path-specific,", meaning that it will depend on the order in which risk factors are assumed to be removed from the population. The SAF for removing hypertension as a risk factor for coronary artery disease will be different if hypertension is removed before smoking, vs. if smoking is removed before hypertension.

Potential solutions to double-counting in the VIMC

As previously mentioned, double-counting of deaths averted in the VIMC arises due to the use of pathogen-specific models, each of which contains the implicit counterfactual assumption that all other Gavi-supported vaccines are provided as actually occurred. In theory, a different modeling approach could be used to avoid this assumption and thus eliminate the problem of double-counting. Rather than running separate models for each pathogen, one unified model could track pathogen-specific death for all pathogens, using a true "no Gavi-supported vaccination" counterfactual and the combined presence of all Gavi-supported vaccinations as the scenario of interest. In practice, this solution is not likely to be feasible given the state of scientific knowledge and the structure of the VIMC.

A separate approach, suggested by the PAF literature, would be to mimic the process of sequential calculation of attributable fractions. For the VIMC, this would involve first determining a "path", that is, a priority ordering of vaccines. Hypothetically, this might mean prioritizing measles vaccine first, then hepatitis B vaccine, then yellow fever vaccine, and so on. The pathogen-specific models would be run in sequence, and each model would output demography as well as expected pathogen-specific cases and deaths. The demography outputs from one model would then be used as demographic inputs for the next model in the priority list. In the hypothetical example, demographic outputs from measles models would be used as inputs for hepatitis B vaccine models, and then demographic outputs from hepatitis B vaccine models would be used as inputs for yellow fever vaccine models. This would eliminate the double-counting problem, because deaths in persons with more than one possible fatal infection would not show up as averted until all the necessary vaccinations were in place. Unlike the creation of a combined model, this approach is technically possible within the VIMC as currently structured. But it has numerous logistical challenges that may make it unrealistic to implement. One particular challenge would be that some vaccines and their associated models (e.g. HPV) are sex-specific. Implementing the PAF-style approach would mean either running all models prior to the sex-specific models using sex-specific demography, or making assumptions about the sex distribution of population outputs from the previous models.

A third approach, also suggested by the PAF literature, is to estimate the total proportion of deaths averted based on the complements of the pathogen-specific deaths averted. To do this,

we would define the "fraction of deaths averted" by vaccinating against a given pathogen *p* (DAF_p) as the number of pathogen-specific deaths averted divided by the total number of deaths (from any cause). The fraction of all deaths averted by all Gavi-supplied vaccines would then be approximated by:

$$DAF_{total} = 1 - \sum_{p} \left(1 - DAF_p\right)$$

(and multiplying DAF_{total} by the total number of all-cause deaths would yield the total number of deaths averted by Gavi-sponsored vaccines). As with the corresponding approach to combining PAF estimates, this approach assumes pathogen-specific deaths are biologically and statistically independent, and only asymptotically approaches 1.0. Thus DAF_{total} will probably under-estimate the combined effect of multiple Gavi-supported vaccination programs. However, it may serve as a useful lower bound for combined estimates of deaths averted.

Next steps

The primary next step is to assess the magnitude of possible double-counting of lives saved within the VIMC. To put an upper bound on the potential for double-counting, the VIMC can calculate DAF_{total} across all the modeled pathogens. The difference between deaths averted as a sum across all models vs. DAF_{total} would be the extreme case for double-counting. If double-counting appears to be trivial based on this measure, no further work would be needed.

If calculating DAF_{total} suggests that meaningful double-counting may be occurring, the second step would be to calculate sequential deaths averted (mimicking the SAF) for a subset of models where the overlap in deaths averted may be most significant. Measles, pneumococcal, and Hib would be good candidates, as the bulk of deaths averted from these vaccines are in the same age group (aged <5 years). As an example, for a country with high vaccine impact from all three vaccines (e.g. Nigeria), these three models could be run in sequence, with demography from one model used as inputs to the next model in the sequence. Then, deaths averted would be summed from the chained models and compared to deaths averted from the models run independently. This would provide the actual magnitude of double-counting (under the assumption that the models are chained in the order of programmatic priority).