

I hesitate to offer an option, but here is one way to think about it. I'm supposing that those surveys are representative at the regional level - some aren't, and I don't know about this particular one.

Assuming they are regionally representative, you could normalize first within each region so that the region sums to 1. Then you could use outside data on regional populations to overlay population weighting on the within-region probability weights, doing that separately for each survey round.

The thinking is that even if the weights for some region sum to a number you aren't interested in, the relative size of the weights within a region contains all of the probability of selection information (within the region information). So by forcing those weights to sum to one, you preserve the probability weighting and lose the problem of the weights summing to the wrong number. Then, you multiply all those weights by the population of the region so that the region itself has weight summing to its population. It would look something like this in Stata:

```
*region total sum of weights
egen region_tot_weight = total(weight), by(region)

*re-normalize within region
gen region_norm_weight = weight/region_tot_weight

*overlay population weight
gen final_weight = region_norm_weight * region_population
```

I've run this by some people who should know, and they generally seem to think it makes sense, but I wouldn't say that this is guaranteed right. I've never seen it used in a paper. Then again, most papers don't really address the weighting problem or at least don't provide any information on how they actually adjusted the weights. I think the widespread use of multiple-round DHS analysis is pretty recent and demand is a bit ahead of the technical expertise in this area.