Abstract: The attributable fraction (AF) is a widely used measure to assess the impact of an exposure on a disease. It is commonly estimated through maximum likelihood, which requires a regression model for the outcome. Recently, it was demonstrated that the AF can also be estimated through inverse probability weighting, which requires a model for the exposure. In this presentation, we show how to construct doubly robust estimators for the AF. These estimators require one model for the outcome and one model for the exposure, and are consistent if either model is correct, not necessarily both. Thus, a researcher will have two chances, instead of only one, to make valid inference. We consider both cohort/cross-sectional studies and case–control studies.