

Average Treatment Effect versus Effect of Treatment on the Treated: Some Intuition.

During gestation of a high-risk pregnancy, two clinical options are available to manage the risk of death: induction of premature delivery and expectant management. Suppose a researcher is interested in quantifying the effect of inducing delivery prematurely on fetal and infant death. This researcher collects data on a cohort of high-risk pregnant women, including whether delivery was induced prematurely, fetal/infant death, and a host of confounding variables. All parties involved agree the study is designed perfectly (no confounding, measurement error, loss to follow-up). They calculate the average treatment effect of premature delivery induction on fetal and infant death on the risk difference scale:

$$E(Y^1 - Y^0) = 0.15. \quad (\text{ATE})$$

This researcher concludes that, if all high-risk pregnancies were induced prematurely 15 more out of every 100 would end in death, relative to what would happen if all high-risk pregnancies were left to expectant management. In light of this incredibly high excess risk of death, this researcher advises abandoning the practice of premature delivery induction.

Another researcher questions the relevance of the average treatment effect. They argue that physicians would never induce delivery prematurely in all versus no high-risk pregnancies. Rather, the more interesting question is: for those women whose pregnancies were actually induced, what would the risk of death have been had they not been induced? This researcher thus calculates the effect of treatment on the treated on the risk difference scale:

$$E(Y^1 - Y^0 \mid A = 1) = -0.05. \quad (\text{ETT})$$

This other researcher concludes that, among those whose pregnancies were actually delivered prematurely, the risk of death would have been higher had they not been delivered prematurely.

This example demonstrates the fundamental difference between the ATE and the ETT: for those high-risk pregnancies that *were not* induced prematurely, the act of inducing premature delivery would not be beneficial. But for those high-risk pregnancies that *were* induced prematurely, the act of inducing premature delivery was beneficial. The ATE averages the beneficial and non-beneficial effects in the entire population, to give an overall non-beneficial effect. The ETT isolates the beneficial effect among those who actually received the intervention. Thus, in this example, premature delivery actually did benefit those who received it, even though it would not benefit everybody.