Bayesian estimators ~~and design optimization~~ for play-the-winner designs

- Introducing Play-the-Winner
- Problems with unbiased (frequentist) estimators
- The Bayesian play-the-winner model
- Other applications of cherry picking correction: screening for biomakerk pharmaceutical target

Helene H Thygesen
Odense University Hospital
Odense, Denmark

Vincent van der Noort
Netherlands Cancer Institu
Amsterdam, Netherland

Play the winner - the design

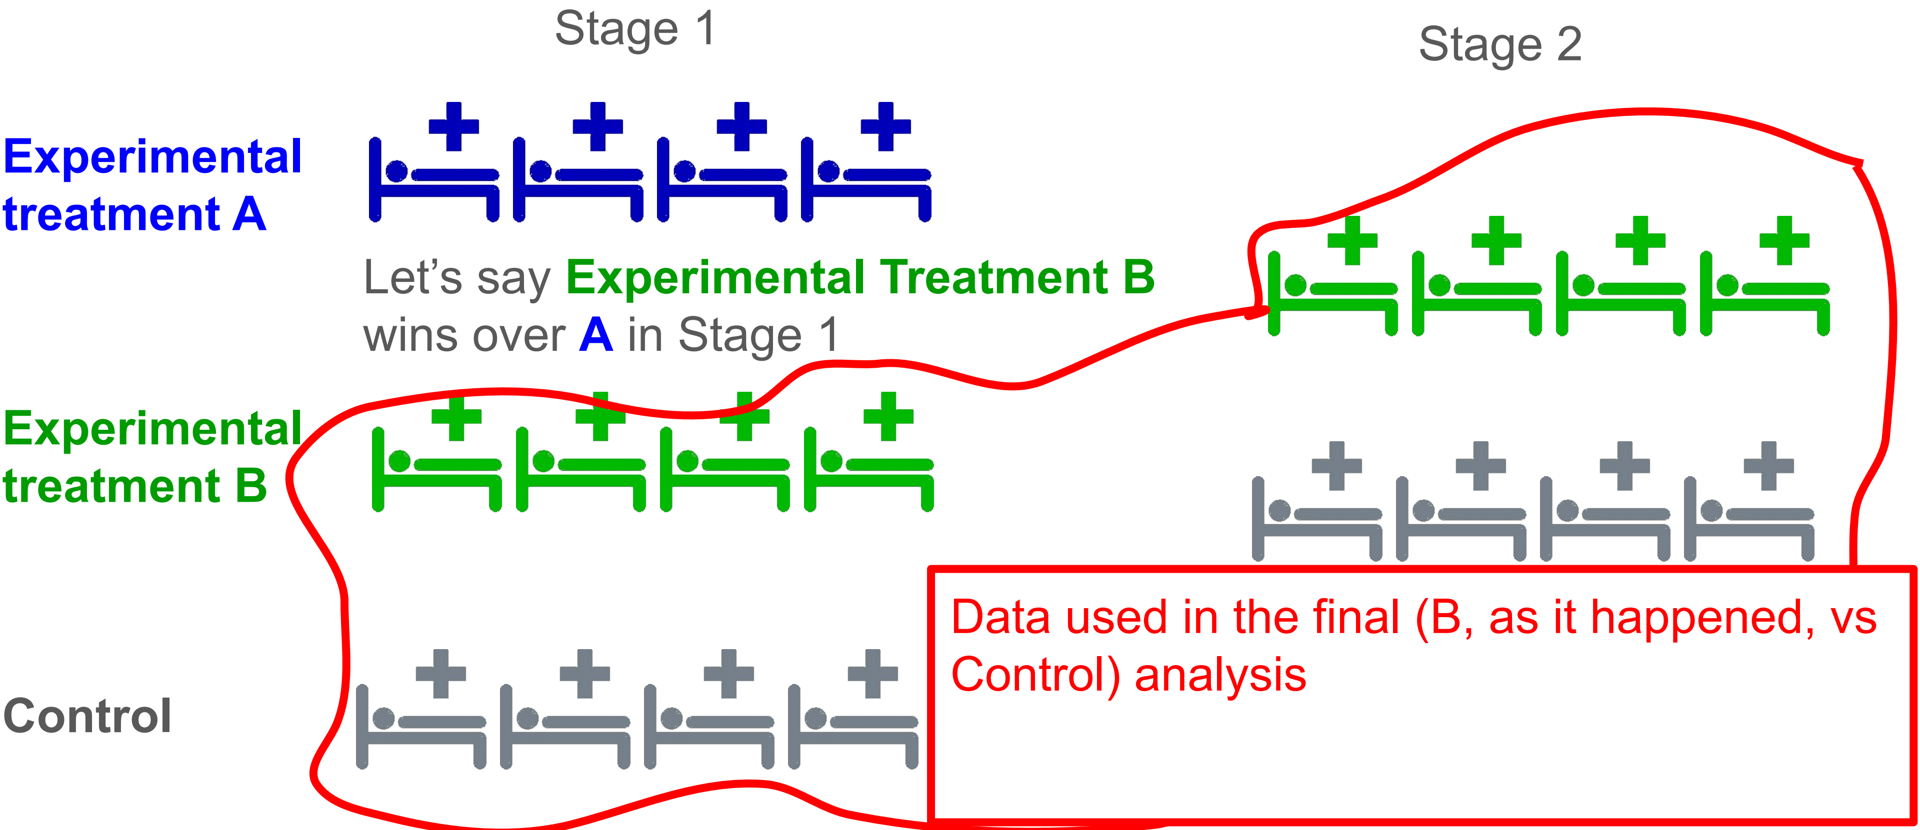

Play the winner - the model and estimand

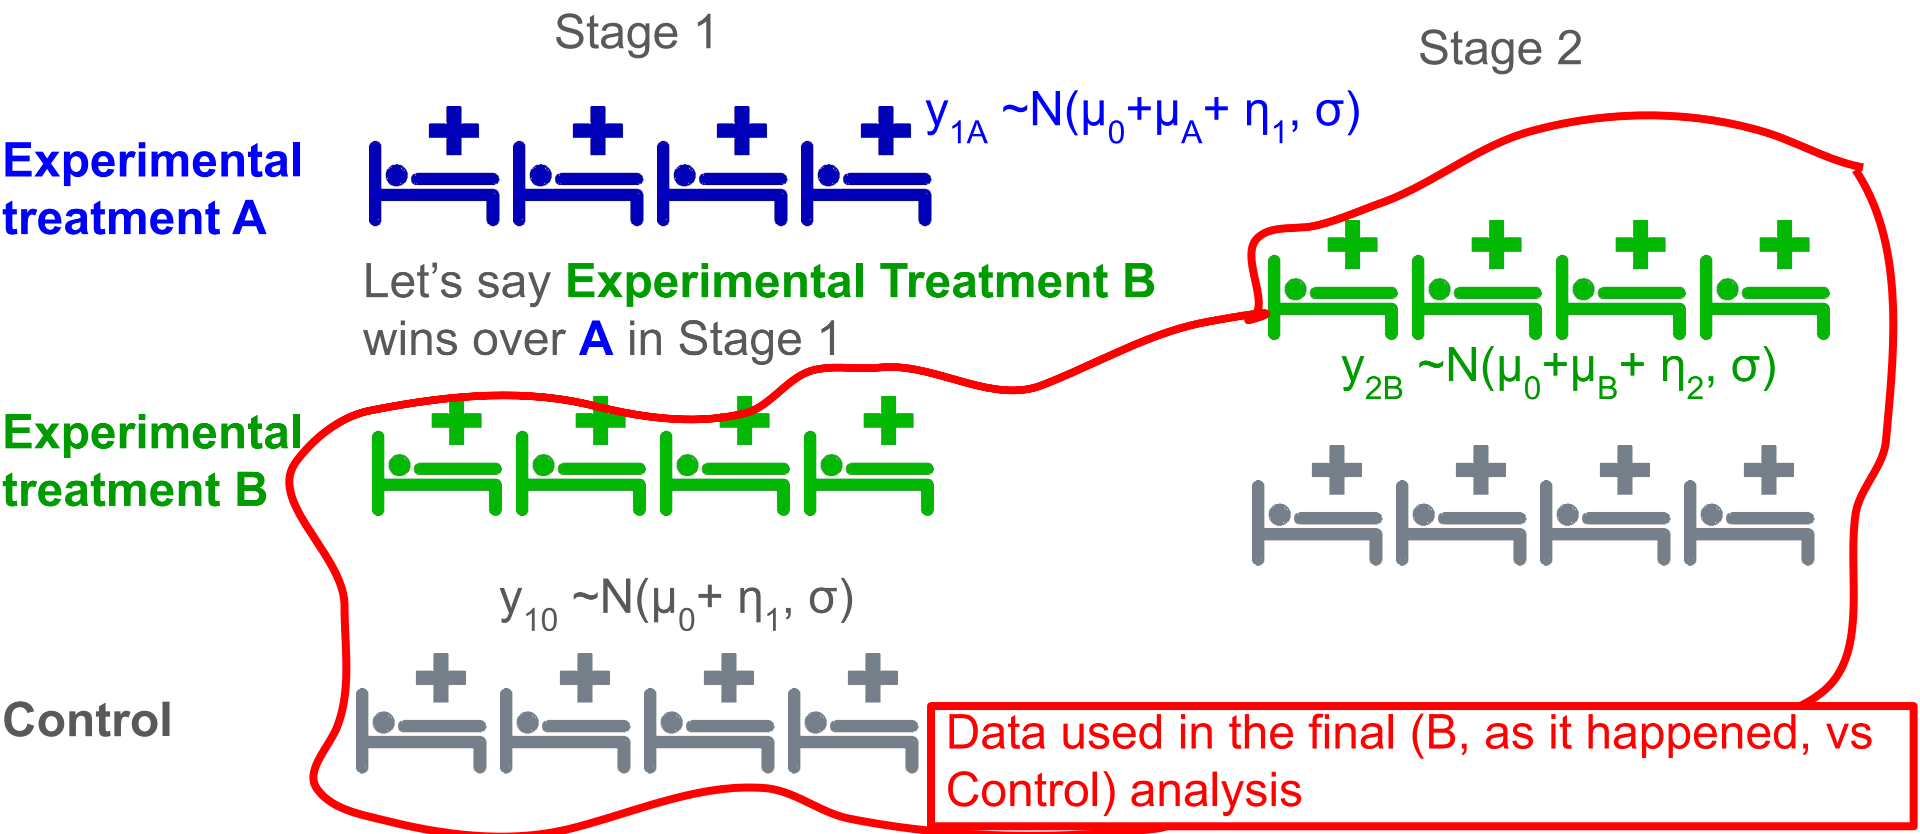

The frequentist analysis: it is difficult to avoid bias on the estimate of μ_{OPT} (index OPT defined as the treatment that appeared to be superior)

Suppose $\mu_B = \mu_A = 0$. The naive estimator $\hat{\mu}_B = m_B - m_0$ is biased because, by design $m_B > m_A$ so it is plausible that is also $m_B > 0$.

$m_B - m_A > 0$ is biased for $\mu_B - \mu_A = 0$.

$m_B - m_A = 0$ is biased for $\mu_B - \mu_A > 0$.

Posterior mean of μ_{OPT} (say, OPT = B)

If $\mu_B > \mu_A$, Epost (μ_{OPT}) = m_B

If $\mu_B < \mu_A$, Epost (μ_{OPT}) = m_A

So:

Epost (μ_{OPT}) = post($\mu_B > \mu_A$) m_B + post($\mu_B < \mu_A$) m_A

Every effect size estimator is biased in the frequentist sense

An estimator $\hat{\theta}$ for **θ** is said to be unbiased if for every hypothetical true **θ**, the expected bias is zero

$\forall$ **θ**: E($\hat{\theta}$ -**θ**|**θ**) = 0

Suppose **θ** = 0. Since $\hat{\theta} \geq 0$ by design, the only unbiased estimator is $\hat{\theta}$ = 0. Which is biased for every other value of **θ** !

Beyond play-the winner

Suppose you want to develop a new biomarker (or therapeutic target)

You find umpteen candidates that have been studied and reported in the litterature

You cherry-pick the one that reportedly worked best!

Q:

What effect size do you assume for you sample size calculations etc?