Bayesian Safety Surveillance with Bias Correction

Fan Bu (UCLA)

on behalf of the OHDSI team collaborating with the FDA BEST Initiative

Detecting risks of adverse events post vaccination

H_0: no increased risk (no signal) v.s. H_1: increased risk (signal)

- Key challenge: sequential data, with monthly/quarterly updates
- Standard practice: MaxSPRT (sequential test) + Historical Comparator (design)

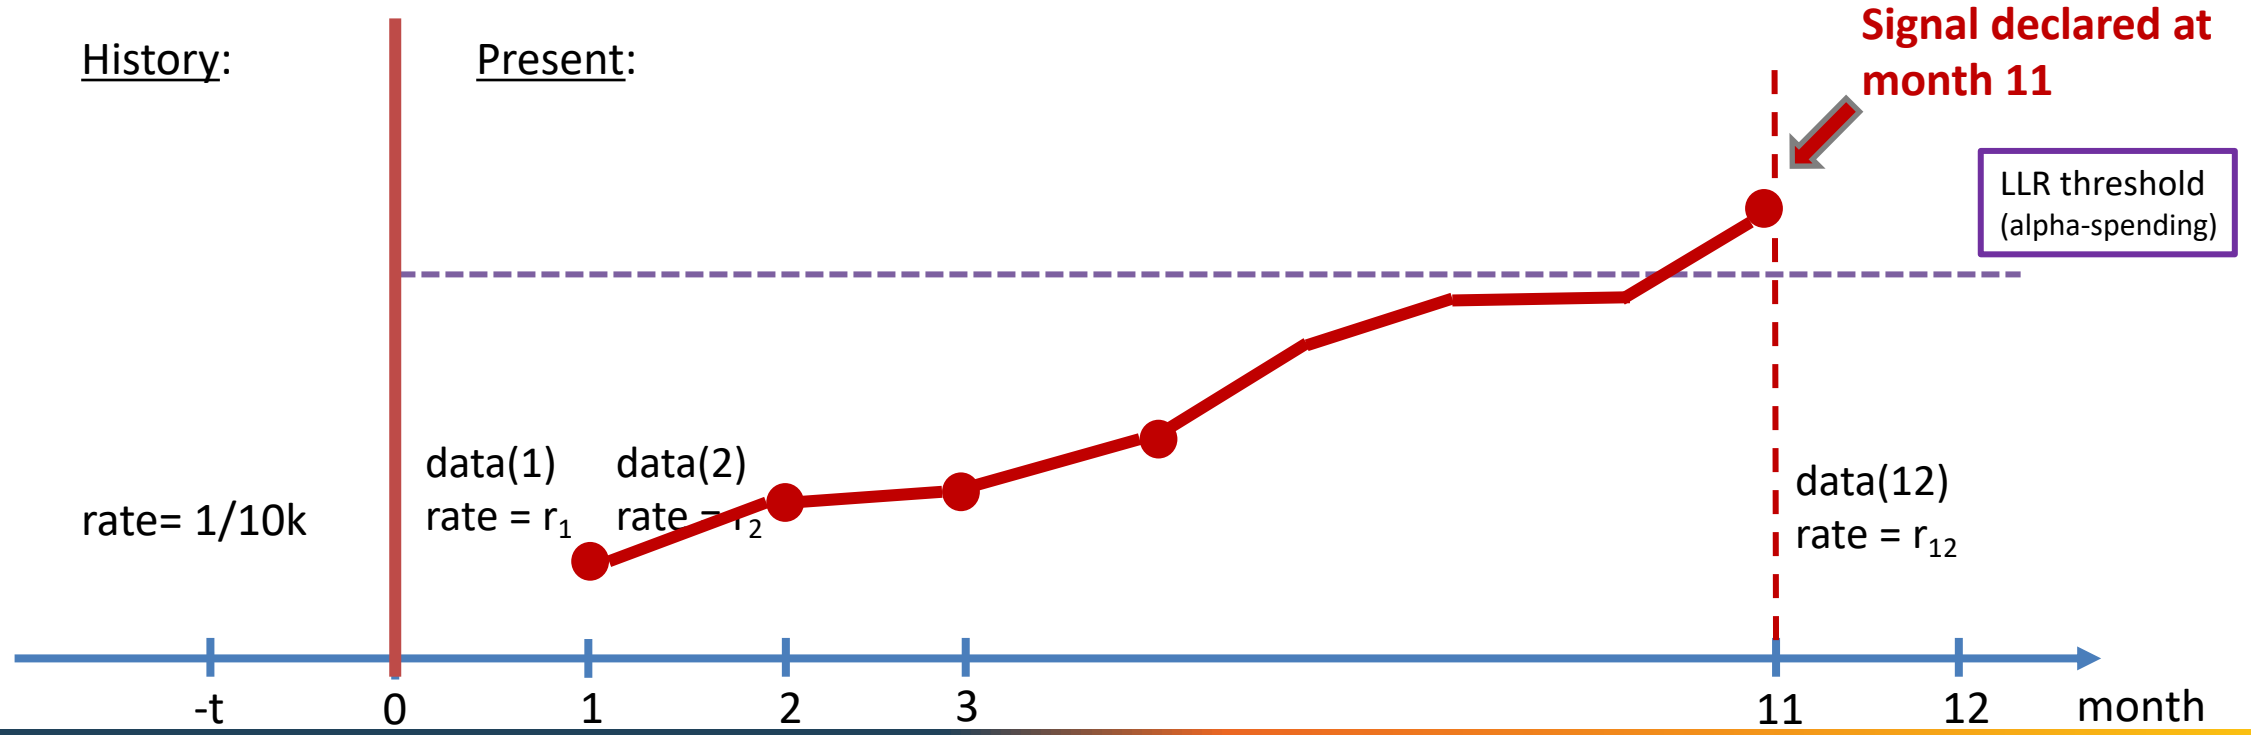

MaxSPRT comes with burdens

- A **fixed** schedule must be pre-specified (to compute threshold)

Month	1	2	……	12
Sample size (incremental)	100	150	……	125

- What if …
  - want to run longer? (e.g., 15 months, 2 years)
  - want to analyze at more frequent or coarse intervals?
  - have different sample sizes in real data?
- Not easy to correct for bias (systematic error)

Type 1 error inflation

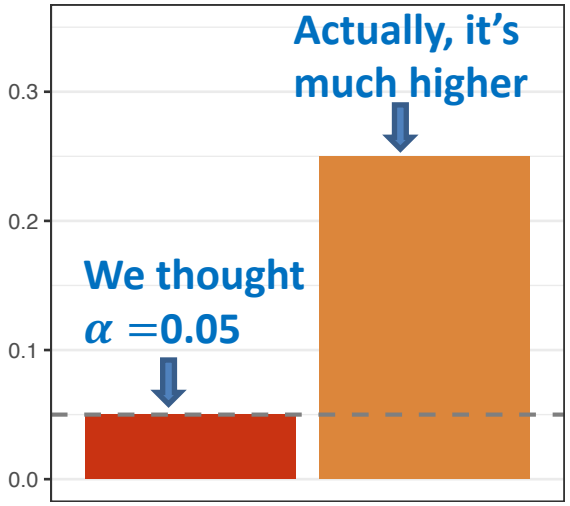

We wish to develop a better alternative

- Bayesian sequential testing framework (details later)
- Comprehensive evaluation on real administrative databases:
  - Retrospective analysis of monthly data
  - Exposures: past vaccines (influenza, H1N1, HPV, Zoster etc.)
  - Outcomes: negative & positive control outcomes + one special outcome of interest (Guillain-Barré syndrome)
  - Evaluation of testing errors, time-to-detection, estimation accuracy, etc.
  - Benchmark against MaxSPRT (with historical rates design)

A Bayesian sequential analysis alternative

Testing via posterior probability $P(H_1|data)$ using Bayesian posterior inference

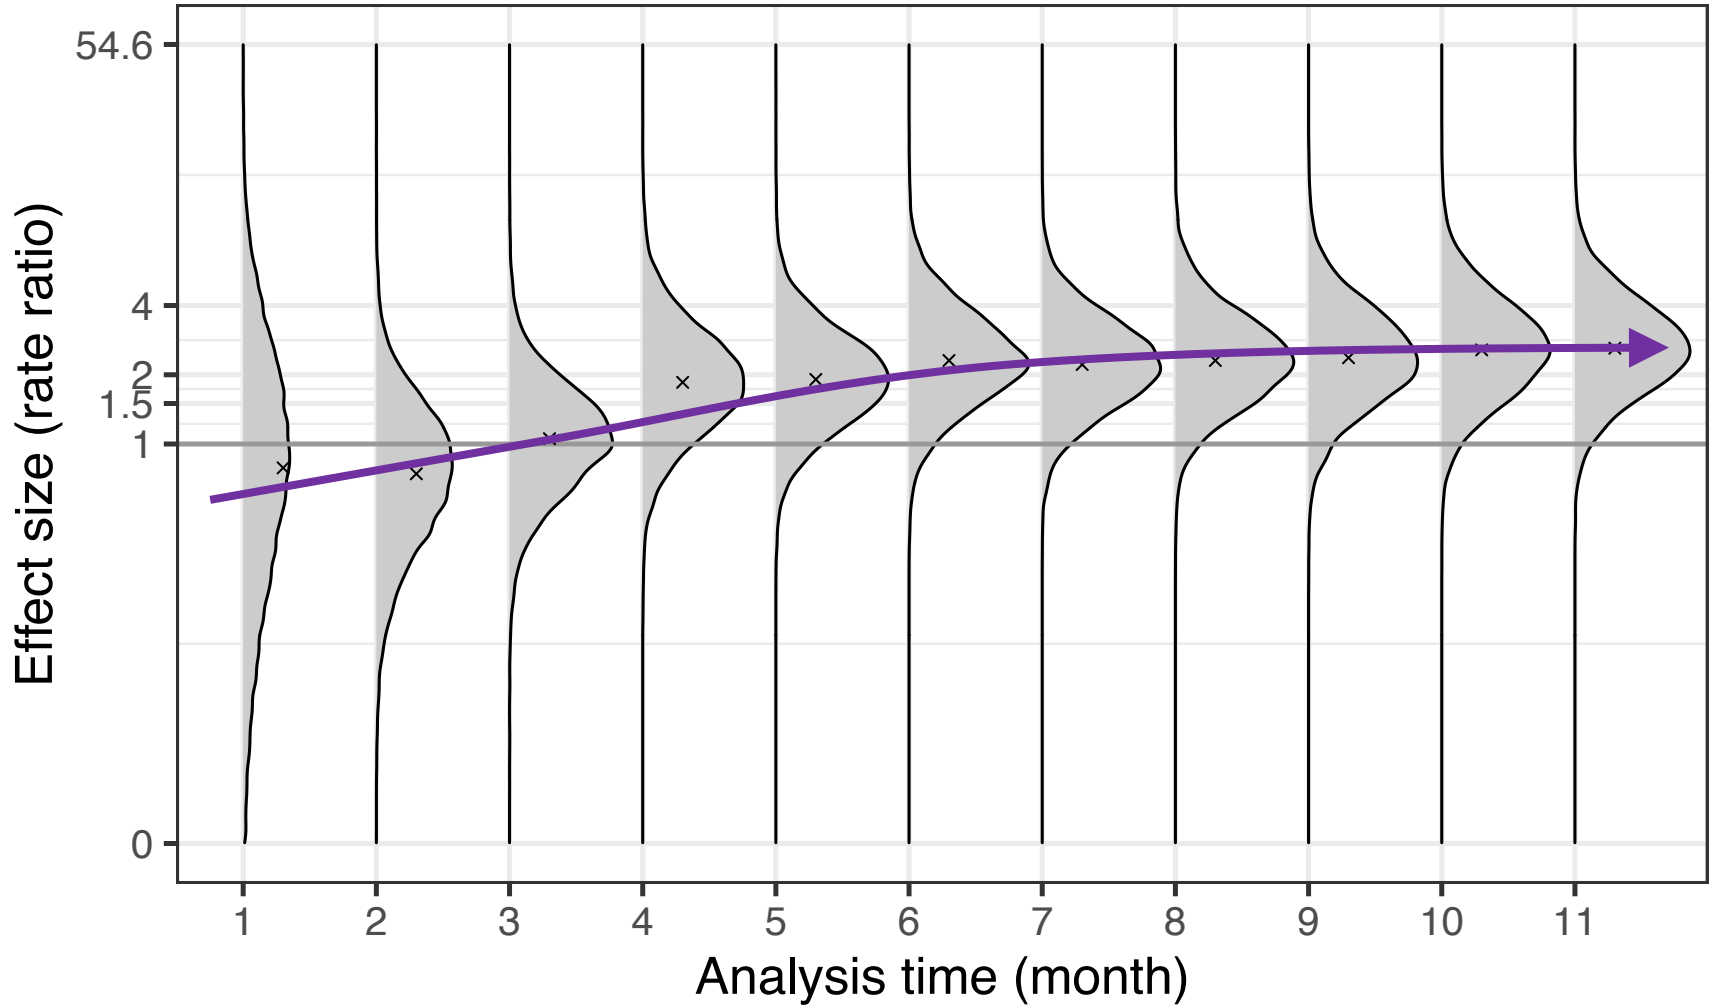

- **Update data evidence via posterior distribution**
- **Retains rich information**
- **Flat: weaker evidence**
- **Concentrated: stronger evidence**

Posterior probability statements from dynamic updates

Testing via posterior probability $P(H_1|data)$ using Bayesian posterior inference

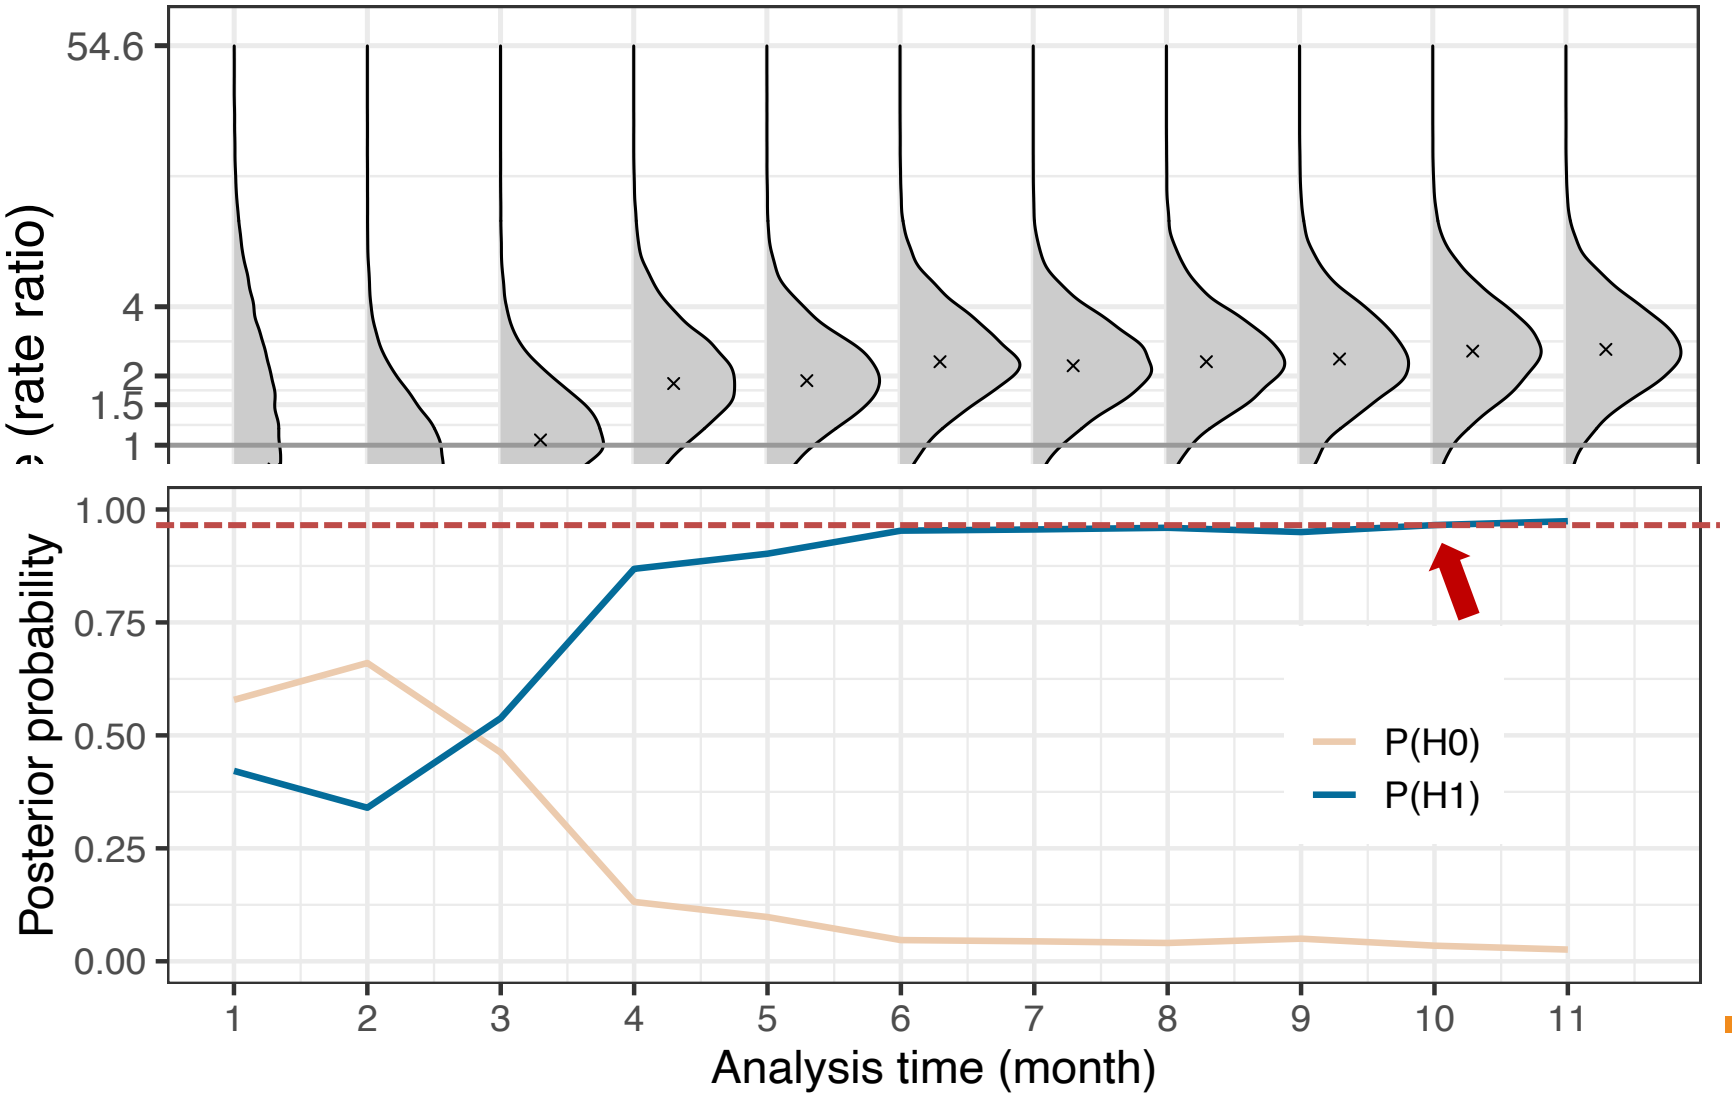

Direct, interpretable statement about hypotheses via

$P(H_i \mid data)$

(not offered by p-value!)

A Bayesian sequential analysis alternative

- Posterior probability $P(H_1|data)$ using Bayesian posterior inference
- Joint model for bias correction using negative control analysis

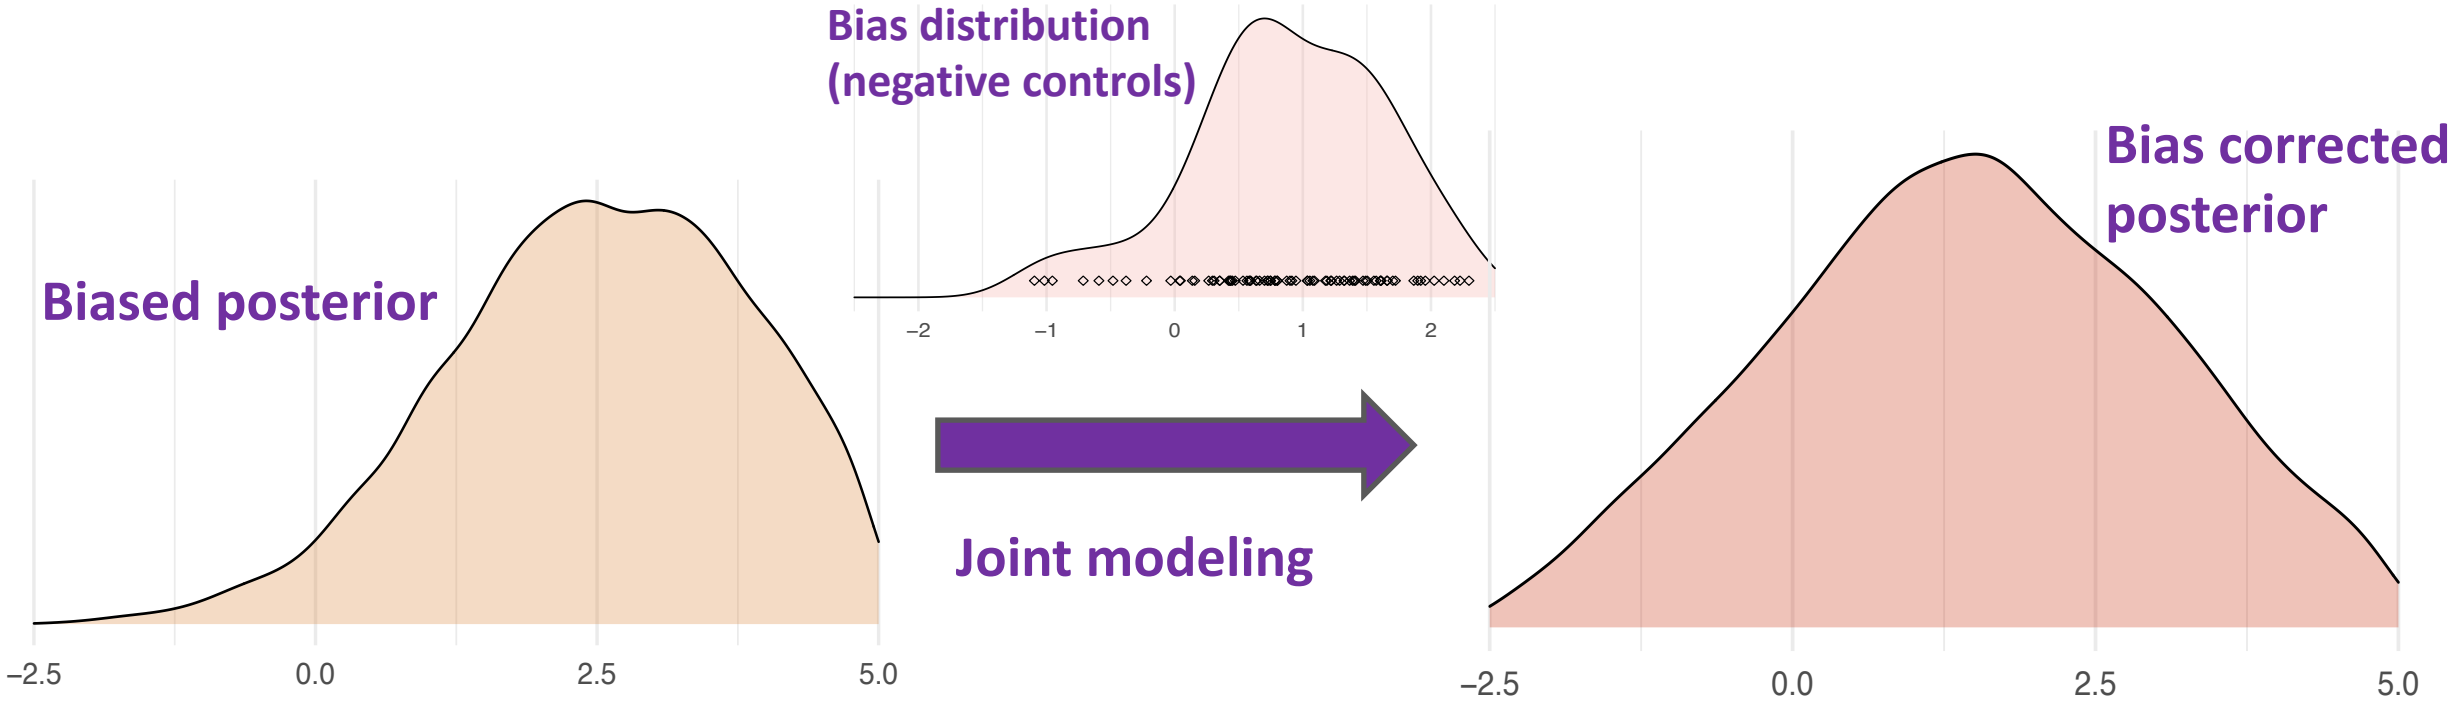

Bayesian bias correction helps control Type 1 error

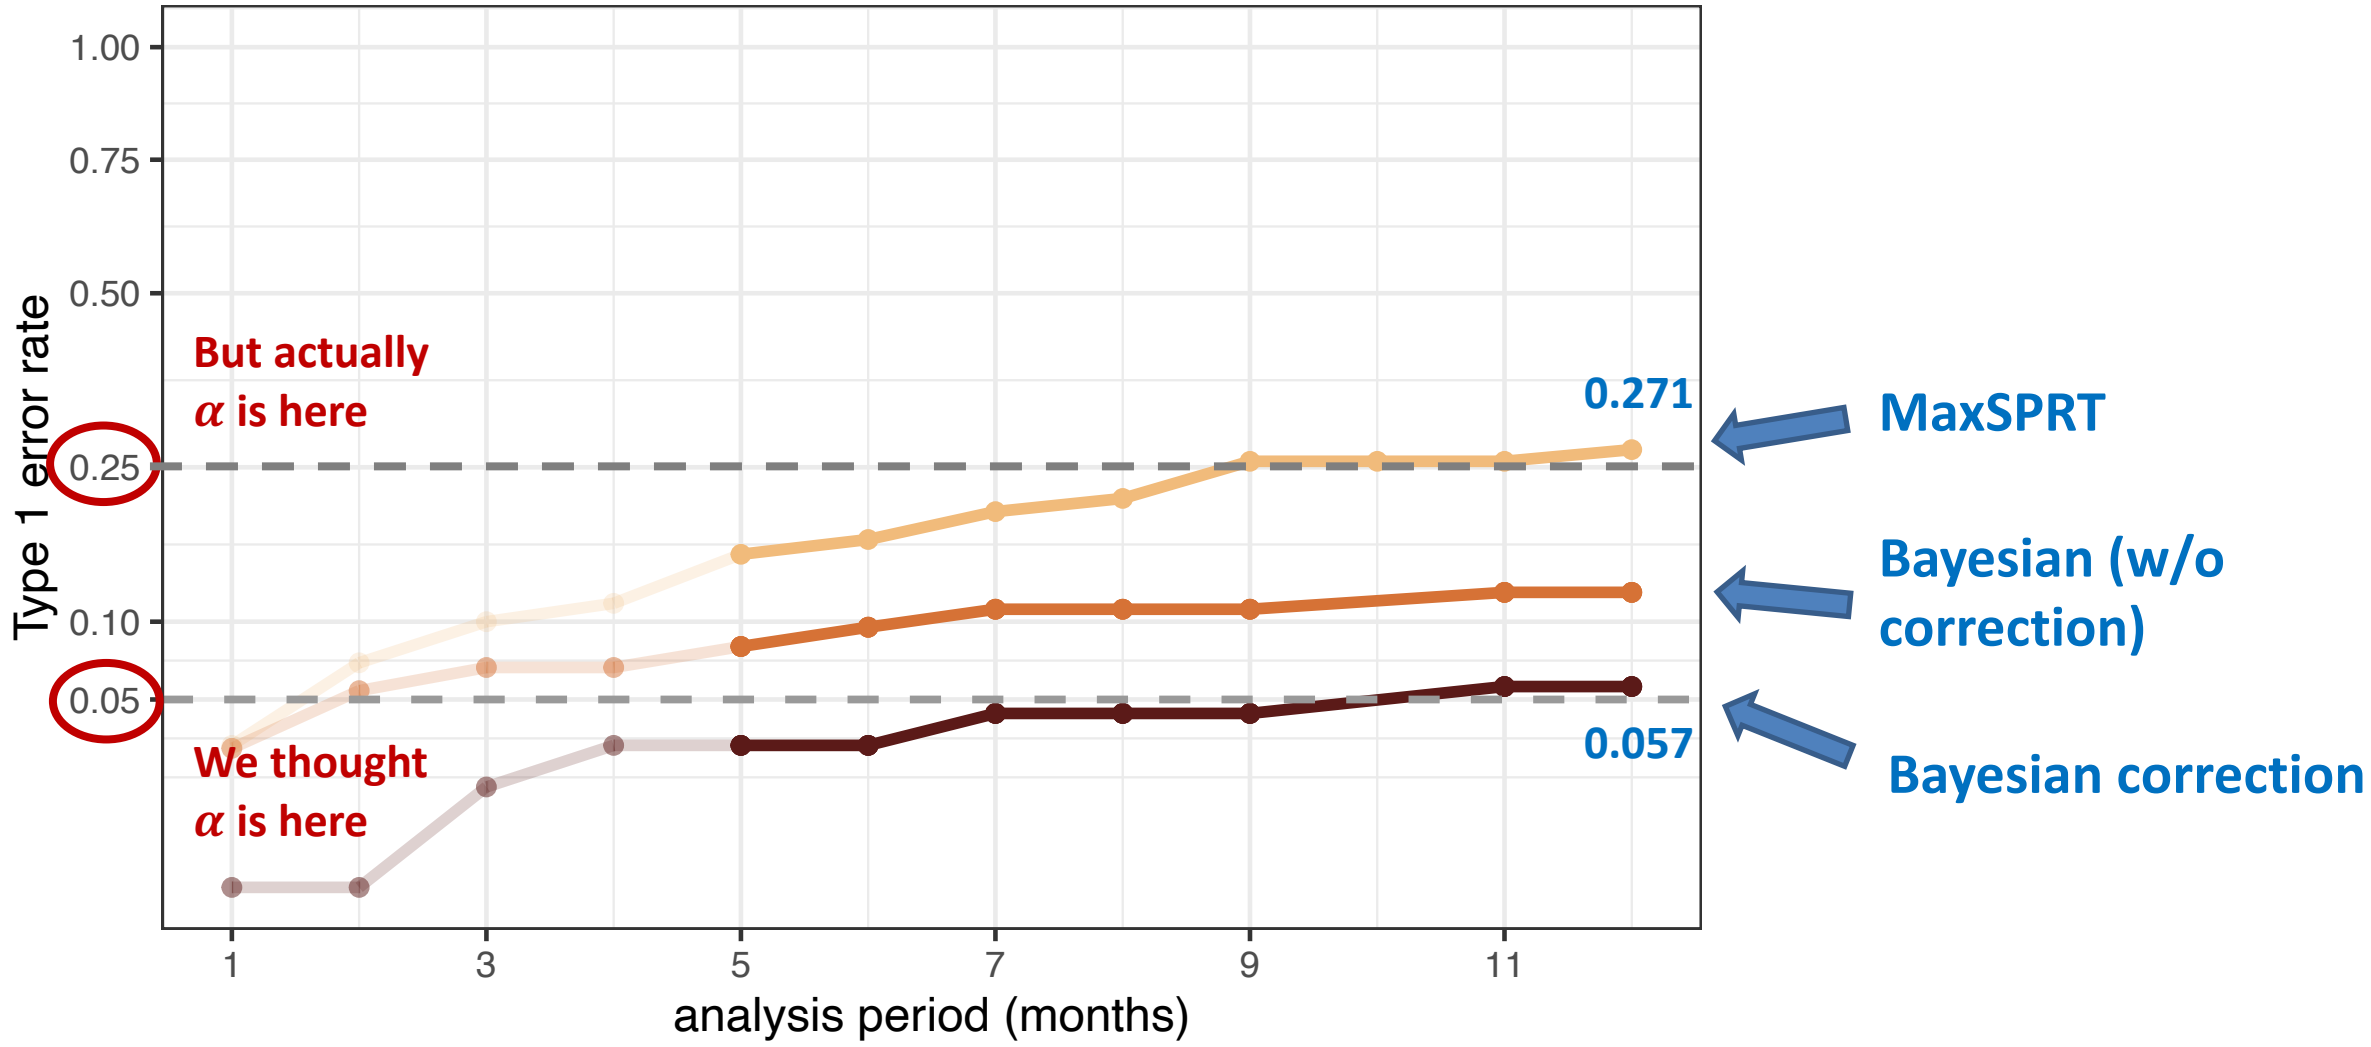

If allow higher α, Bayesian can give higher power

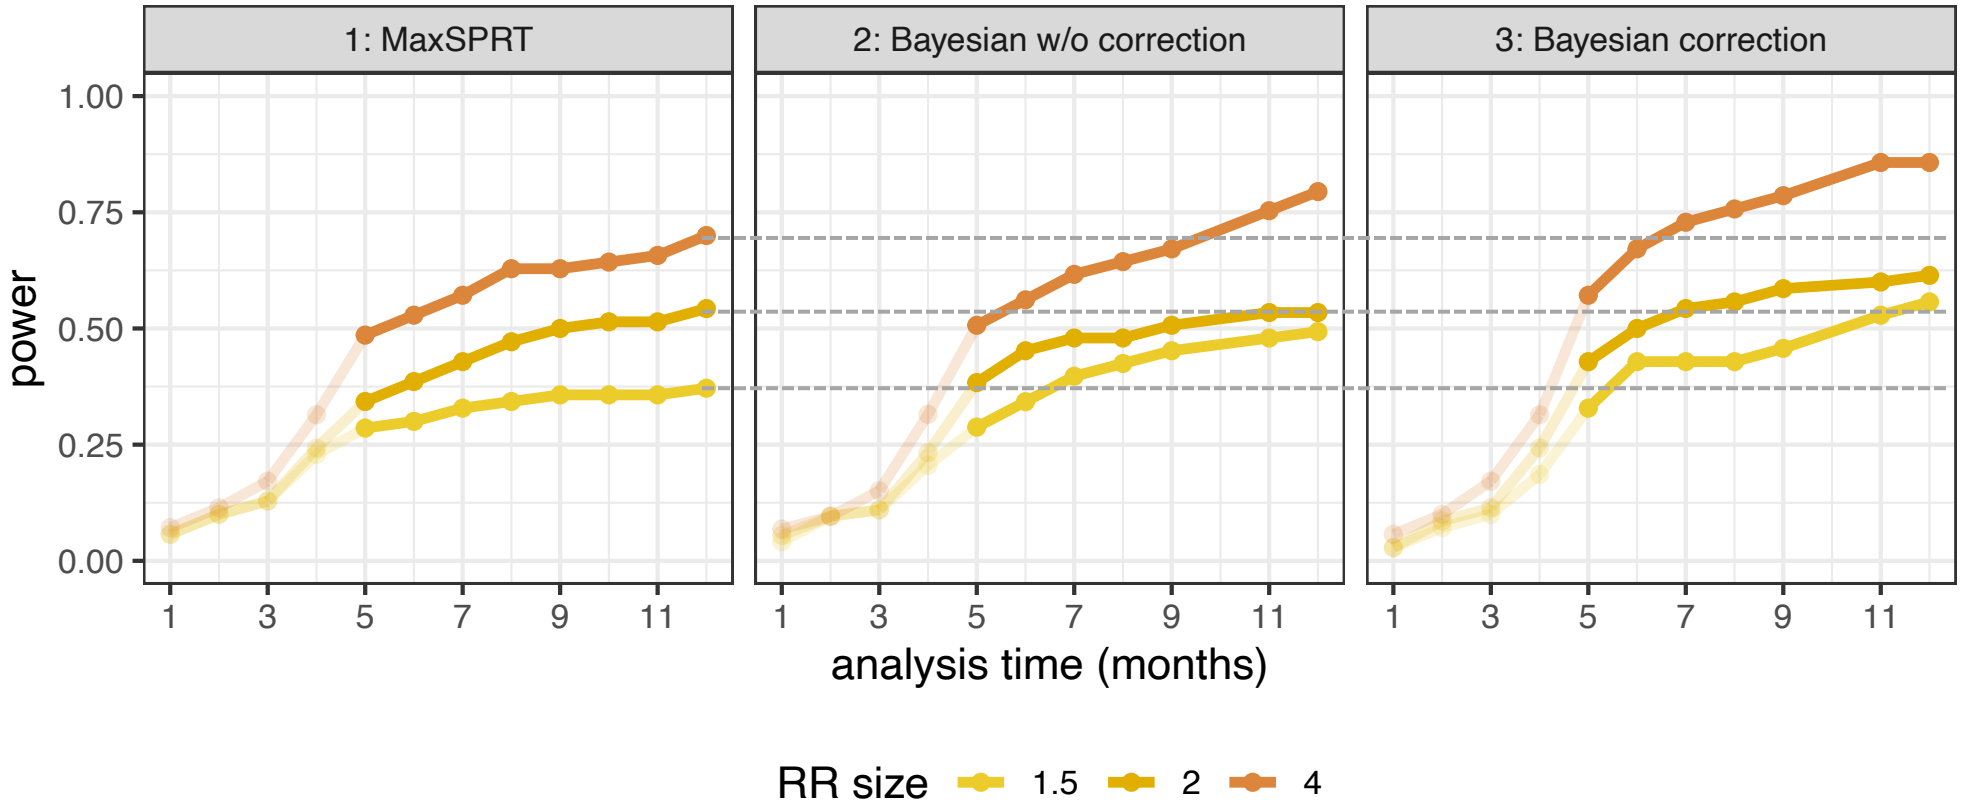

Matching MaxSPRT Type 1 at 0.271: statistical power for different effect sizes.

The tradeoff between α and power

- For initial screening, high-sensitivity but low-specificity might be desirable (don't want to miss a signal!)
- However, claiming (& believing) $\alpha = 0.05$ but in fact $\alpha >> 0.05$ is not best statistical practice
- MaxSPRT handles sequential multiplicity but not bias
- A Bayesian alternative can target bias in a coherent, principled manner
- Why not relax α & use a better method?

Summary

- Bayesian framework presents a better alternative for safety surveillance
  - **Flexible**: no need for fixed schedule
  - **Less bias**: targets bias (from systematic error) & decreases Type 1 error
  - **Powerful**: has sufficient power, esp. if higher α allowed
- Still work to be done:
  - (ongoing) Double threshold testing with futility early-stopping (saves time when confident about safety!)
  - (ongoing) Theoretical/empirical validation of error control with unbounded time horizon
  - Plug-and-play designs without likelihood?

Resources

- Team (@OHDSI):
  - Fan Bu
  - George Hripcsak
  - Kristin Kostka
  - David Madigan
  - Jody-Ann McLeggon
  - Aki Nishimura
  - Patrick Ryan
  - Louisa Smith
  - Marc Suchard
- Special acknowledgements to FDA CBER center for guidance and support!
- Links:
  - Study protocol: **https://suchard-group.github.io/Better/Protocol.html**
  - Study package: **https://suchard-group.github.io/Better/**
- Contact to participate: **Marc Suchard** (Teams) or **msuchard@ohdsi.org**