Optimal Screening Designs with Flexible Cost Structures

Janis Hardwick

Quentin F. Stout

University of Michigan

www.eecs.umich.edu/~jphard

www.eecs.umich.edu/~qstout

Research partially supported by NSF.

April 2002

What are Screening Trials?

Evaluate large number of agents to identify promising ones for further study.

Typically agents are potential medical therapies, but approach also applies to non-medical applications.

Each agent tested independently of all others.

Want average trial to be inexpensive, accurate.

Assume screening test (*pilot study*) has binary response, each agent has an unknown success parameter π.

Given cut-point $\mathbf{C} \in [\mathbf{0}, \mathbf{1}]$, an agent is positive if $\pi \geq \mathbf{C}$.

Use a Bayesian approach: assume prior distribution on π. Beta priors are used in our examples and were used in related work by others, but techniques apply to any priors.

Previous Work

We focus on the work in three papers:

1. Yao, Begg, Livingston (1996, *Biometrics*)
2. Yao and Venkatraman (1998, *Biometrics*)
3. Wang and Leung (1998, *Biometrics*)

These researchers concentrated on designs optimized for fixed type I and type II error rates: F^+ and F^-

1-Stage

Yao, Begg, Livingston (1996): given fixed F^+ and modified F^- (discussed later), determine fixed (1-stage) agent sample size that minimizes total sample size until first promising agent identified.

- Historical data showed that sample sizes used in practice were far too small.
- They noted benefit of early stopping (curtailment).

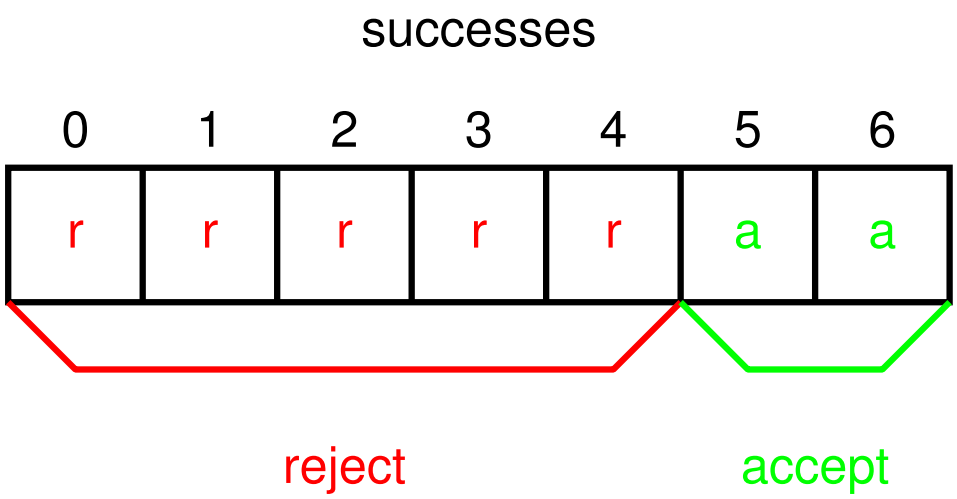

6 observations

2-Stage

Yao and Venkatraman (1998): same constraints and goal, but for 2-stage design. 2nd stage size fixed, but may be omitted (truncation or optional stopping).

- Expected agent and total sample sizes significantly smaller than 1-stage design of Yao, Begg, Livingston.

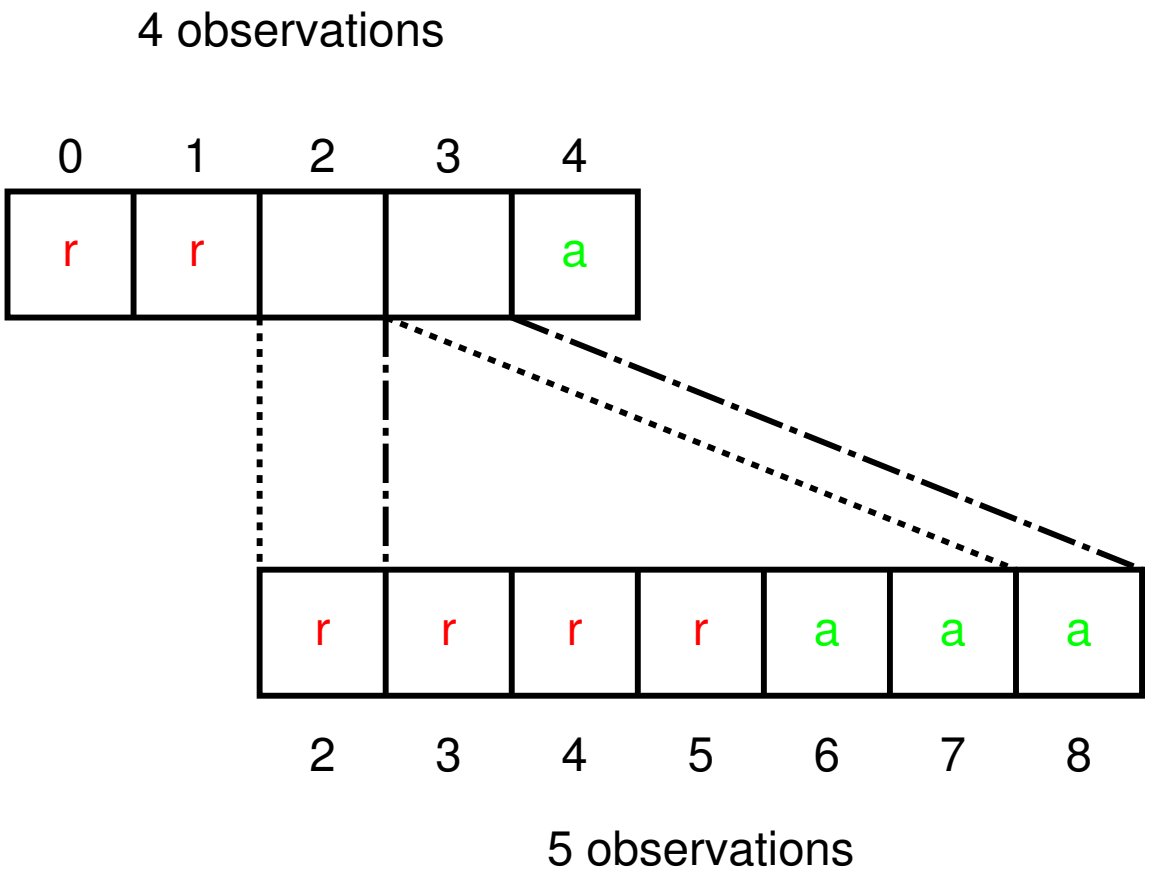

Fully Sequential (∞ stage)

Wang and Leung (1998): minimize total sample size until first promising agent identified, for fully sequential design with optional stopping.

- Expected total sample size minimal, but time maximal.
- Use costs of type I, II errors, not fixed error rates. However, these are not intended to represent true costs, but rather to act like Lagrangian multipliers.
- Optimize total cost = error cost + sample size.
- Unlike the relatively straightforward calculations need for the previous work, here they employ a complex and slow iterative approach, much like a Gittins index calculation.

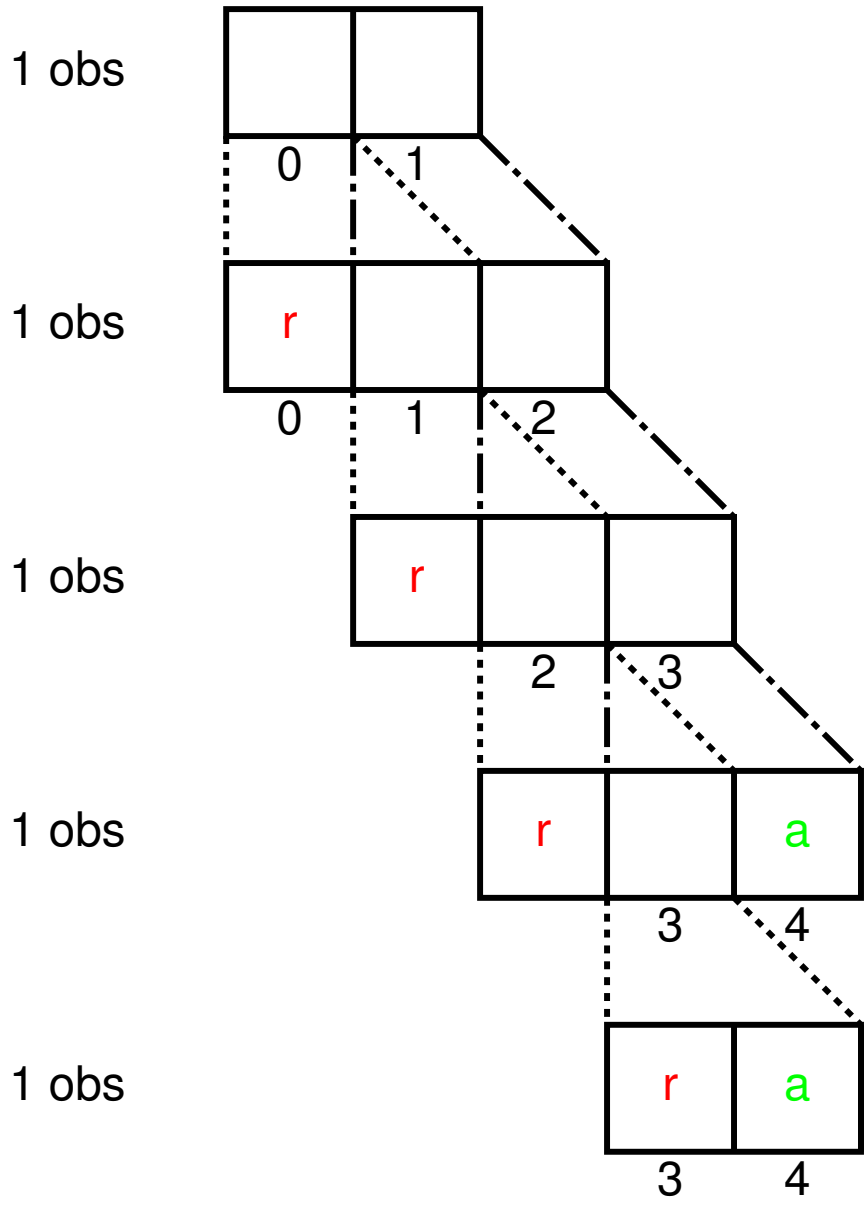
1 obs
0
1
1 obs
r
0
1
2
1 obs
r
2
3
1 obs
r
a
3
4
1 obs
r
a
3
4

Our Approach

Optimize a cost-based model that is realistic, very flexible, and computationally feasible.

Decision-theoretic approach incorporating

- Trial constraints, such as
  - maximum observations per stage
  - maximum number of stages
  - maximum number of observations
- Trial costs, such as
  - setup cost per stage
  - cost per observation
  - cost per failure
- Decision costs, i.e., costs of false positive, false negative decisions. May increase with distance from cut-point.

Goal

Minimize expected total cost per agent

I.e., obeying the trial constraints, minimize the sum of the trial and decision costs.

Note that screening designs are used internally to make proceed/stop decisions, rather than for convincing regulatory agencies that a therapy is efficacious. Thus cost analysis more natural than test of hypotheses.

Computational technique: Dynamic programming.

Example Design

Given prior, cost structure, constraints, trial might be:

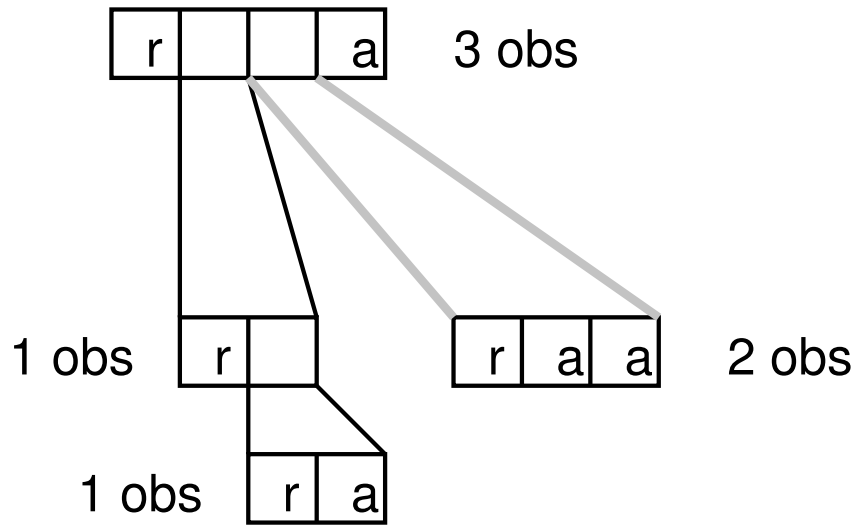

Multistage design. Variable stage sizes,
Variable number stages

Structure determined by costs and constraints, not by fixing F^+ and F^- in advance.

Each step determined by prior and observations.

Optional stopping (truncation) is designed in.

Cost-based approach explicitly incorporates relevant factors.

Previously, one specified false positive and negative rates (F^+ and F^-), trying to take into account the costs of such mistakes versus costs of the screening tests. Typically just a rough guess, especially since cost of screening trial not known.

Making tradeoffs more explicit, and directly optimizing them, improves the decision-making process and quality of the results.

Illustrative Results

- Trial: Unit cost per item, no setup cost per stage.
- Decision: Cut-point = 0.7, FNC (false negative cost) = FPC (false positive cost), use 500, 1000, 2000, 4000.

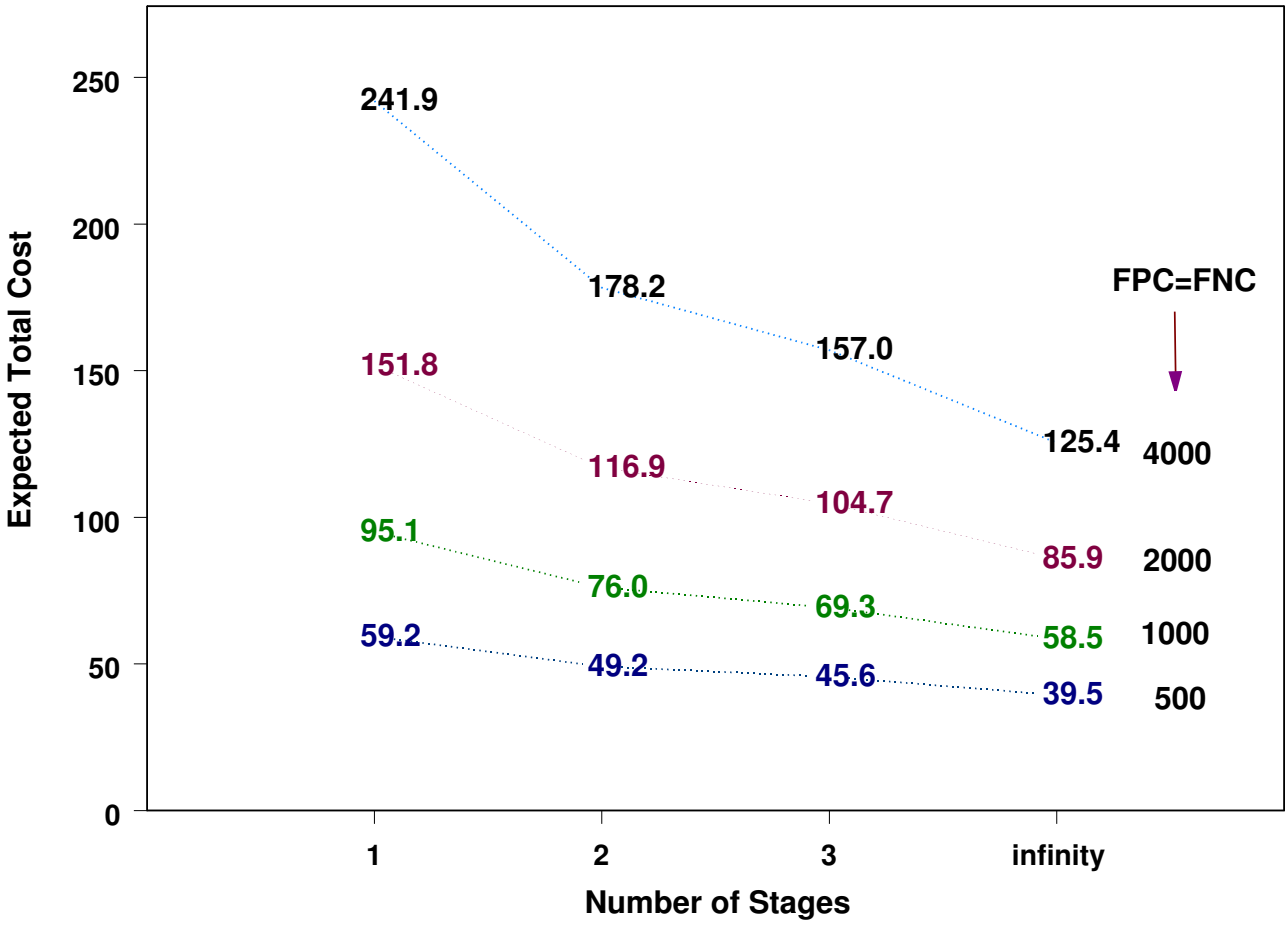

Prior = B(1,1)

Some Bayesian Operating Characteristics

FPC = FNC = 1000

stages	cost	sample size	F^+	F^-
1	95	29	0.110	0.047
2	76	29	0.075	0.036
3	69	28	0.064	0.031
∞	58	26	0.051	0.025

FPC = FNC = 4000

stages	cost	sample size	F^+	F^-
1	242	79	0.068	0.029
2	178	69	0.045	0.020
3	157	66	0.037	0.017
∞	125	56	0.028	0.013

Comparison of 2-Stage Designs

Yao and Venkatraman 2-stage design versus our optimal 2-stage design.

Their requirements:

- False positive rate $F^+ \leq 0.1$
- Prob false negative on any agent until promising agent found ≤ 0.1.

  This is not the same as setting F^- value: if k agents examined until promising agent found, then need $(1\text{-}F^-)^k \geq 0.9$.

We artificially manipulated error costs to achieve their goals.

The Yao and Venkatraman design fixes the size of the second stage, while our design allows it to depend on the outcome of the first stage.

Cut-point	Beta Mean	Opt E(N)	Y & V E(N)	Y & V Excess
0.3	0.2	28.70	35.00	22%
0.3	0.3	24.20	32.80	36%
0.3	0.4	13.46	18.30	36%
0.3	0.5	7.10	7.90	11%
0.6	0.2	65.93	77.10	17%
0.6	0.3	96.35	100.60	4%
0.6	0.4	70.62	81.50	15%
0.6	0.5	37.74	44.50	18%

Prior is Beta with given Mean and Variance = 0.08
(their choices)

Optimization Technique

Optimal design obtained via dynamic programming, working from end of trial towards beginning. All optimization is exact.

For all possible (sample size, successes observed) endings, determine expected false positive and false negative costs. The terminal decision is the one with least cost.

At each intermediate stage, for each (sample size, successes) pair try all options satisfying the constraints, determine which optimizes costs from there to end. This uses costs and decisions computed for the next stage. Note that one option is to stop.

Bayesian framework is critical, allowing one to compute probabilities of outcomes and hence expected cost.

Program efficient, no need for slower iterative computation used by Wang and Leung.

Evaluating Designs

The designs are optimized with respect to given priors, and the operating characteristics shown so far have been determined with respect to these priors.

Some additional evaluations one may desire

- Bayesian: robustness against misspecification of priors
- Frequentist: pointwise determination of costs and F^+, F^- rates.

Exact evaluations are provided for all examples shown, using path induction (Hardwick & Stout 1999).

A wide range of other exact evaluations can be easily performed.

Example Bayesian Robustness Evaluation

Suppose have

- Trial: unit cost per observation, $\leq$ 2 stages
- Decision: FPC = FNC = 1000, cut-point = 0.7

Design Prior Be(1,1)

cost	sample size	F^+	F^-
76	28.7	0.076	0.036

Evaluation Prior Be(3,3)

94	31.9	0.199	0.034

Design Prior Be(3,3)

90	30.4	0.145	0.045

Evaluation Prior Be(1,1)

79	33.2	0.051	0.049

Example Pointwise Operating Characteristics

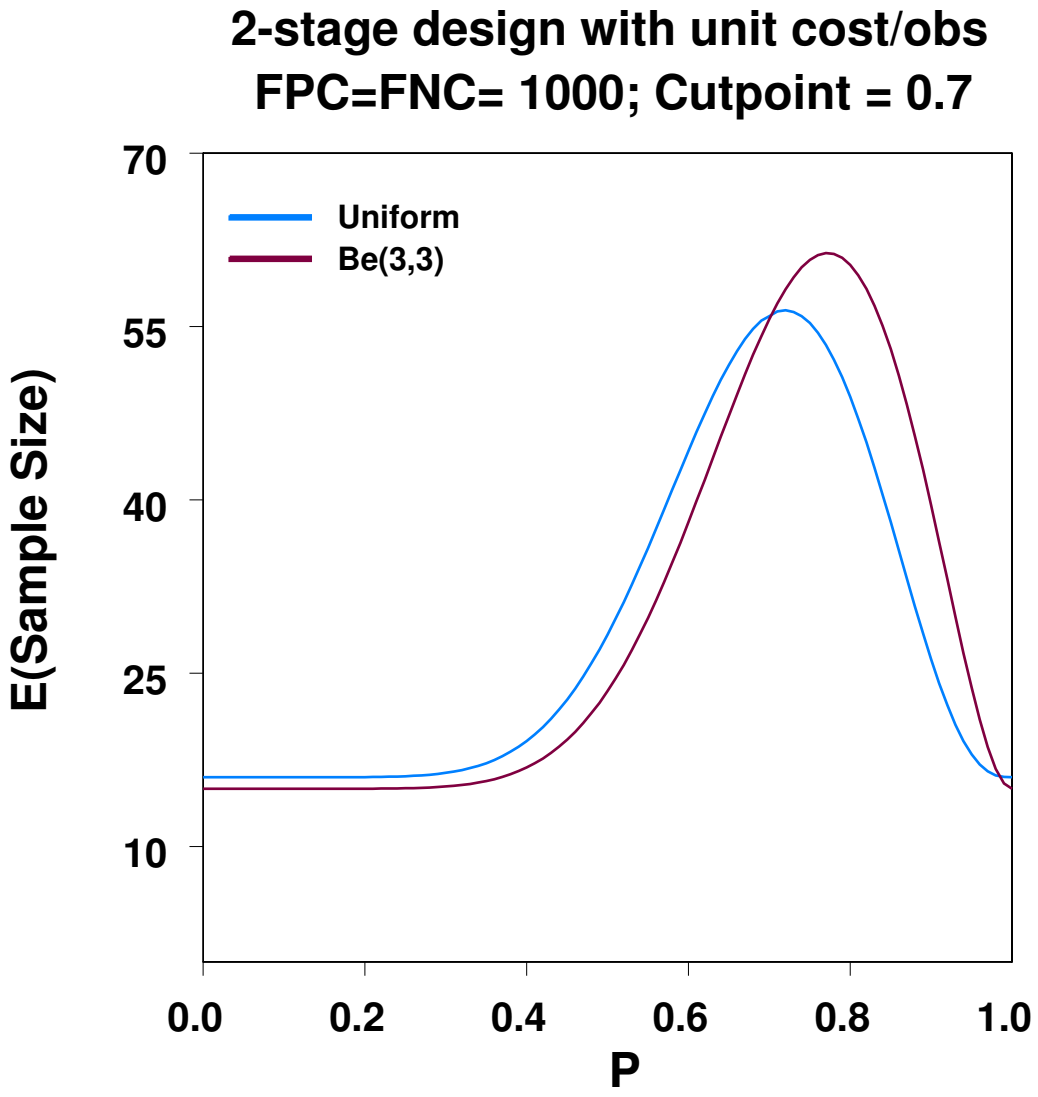

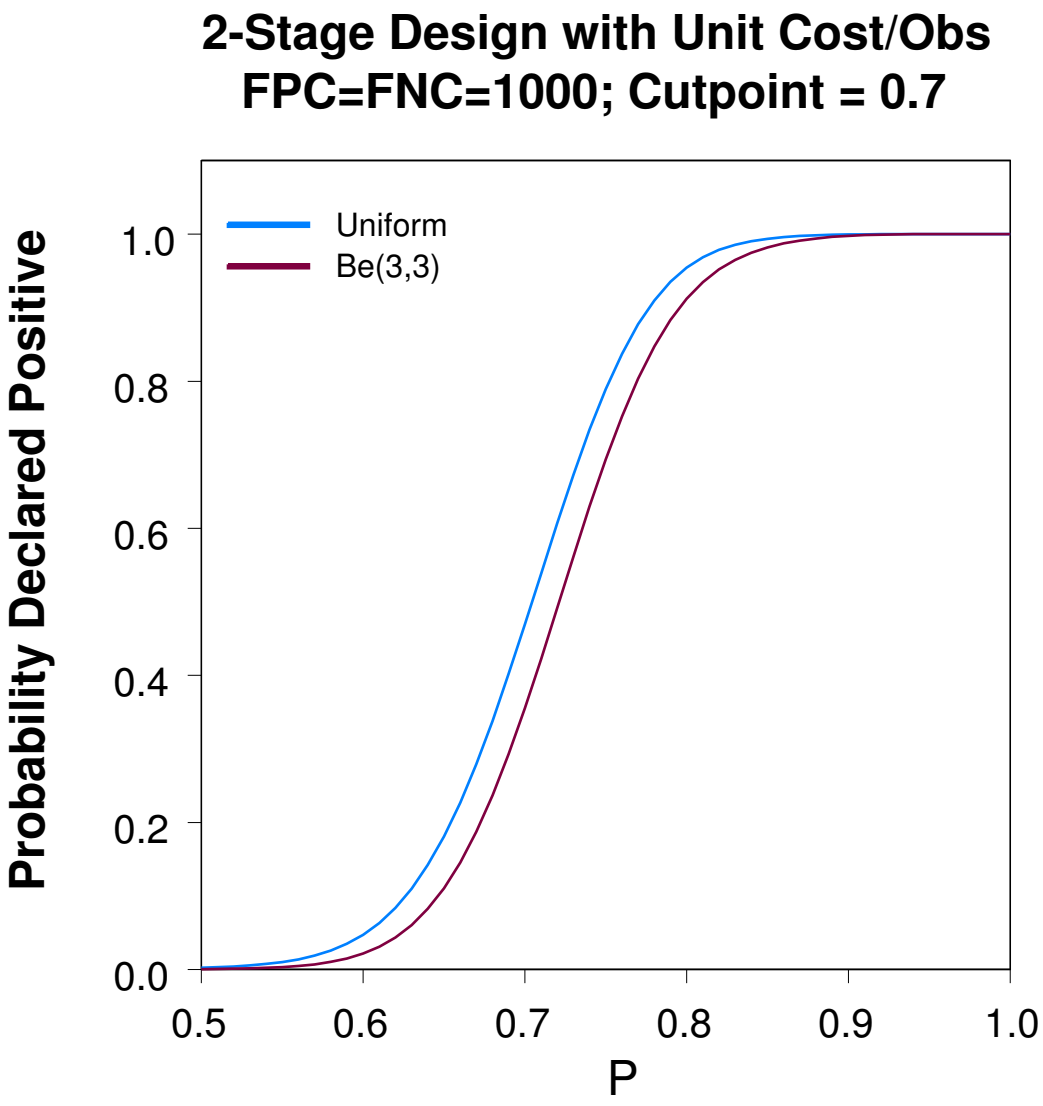
2-Stage Design with Unit Cost/Obs
FPC=FNC=1000; Cutpoint = 0.7
Probability Declared Positive
Uniform
Be(3,3)
1.0
0.8
0.6
0.4
0.2
0.0
0.5
0.6
0.7
0.8
0.9
1.0
P

Reexamining Error Costs

Typically cutoff indicates value where it is expected to be worthwhile to continue.

However, agents below but near cutoff might have payoff if continued, while agents above but near are less likely to have large payoff. I.e.,

> Costs of false positives and false negatives for agents near the cutoff are less than for those far from the cutoff.

Despite this, all previous work, including our examples above, uses step-function costs.

Example Continuous Cost Function

Replace step-function costs with

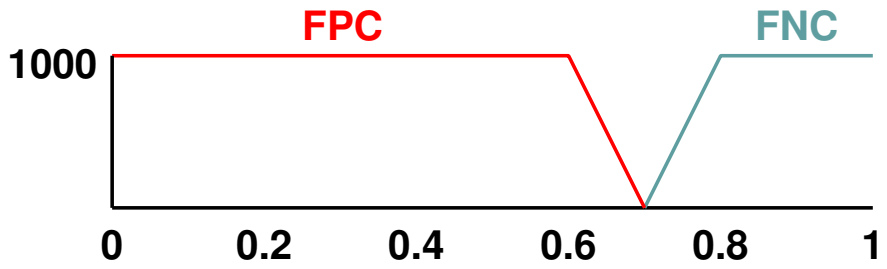

2-Stage Design
FPC=FNC=1000; Cutpoint = 0.7

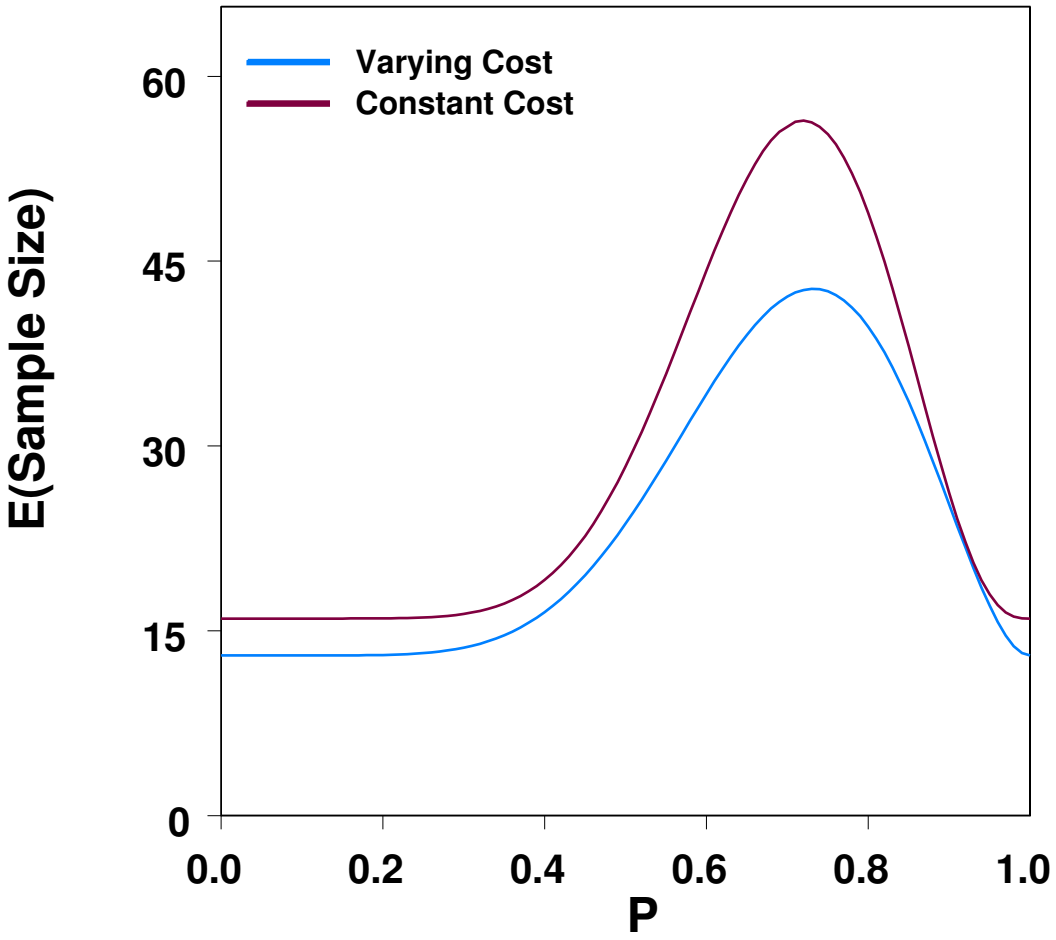

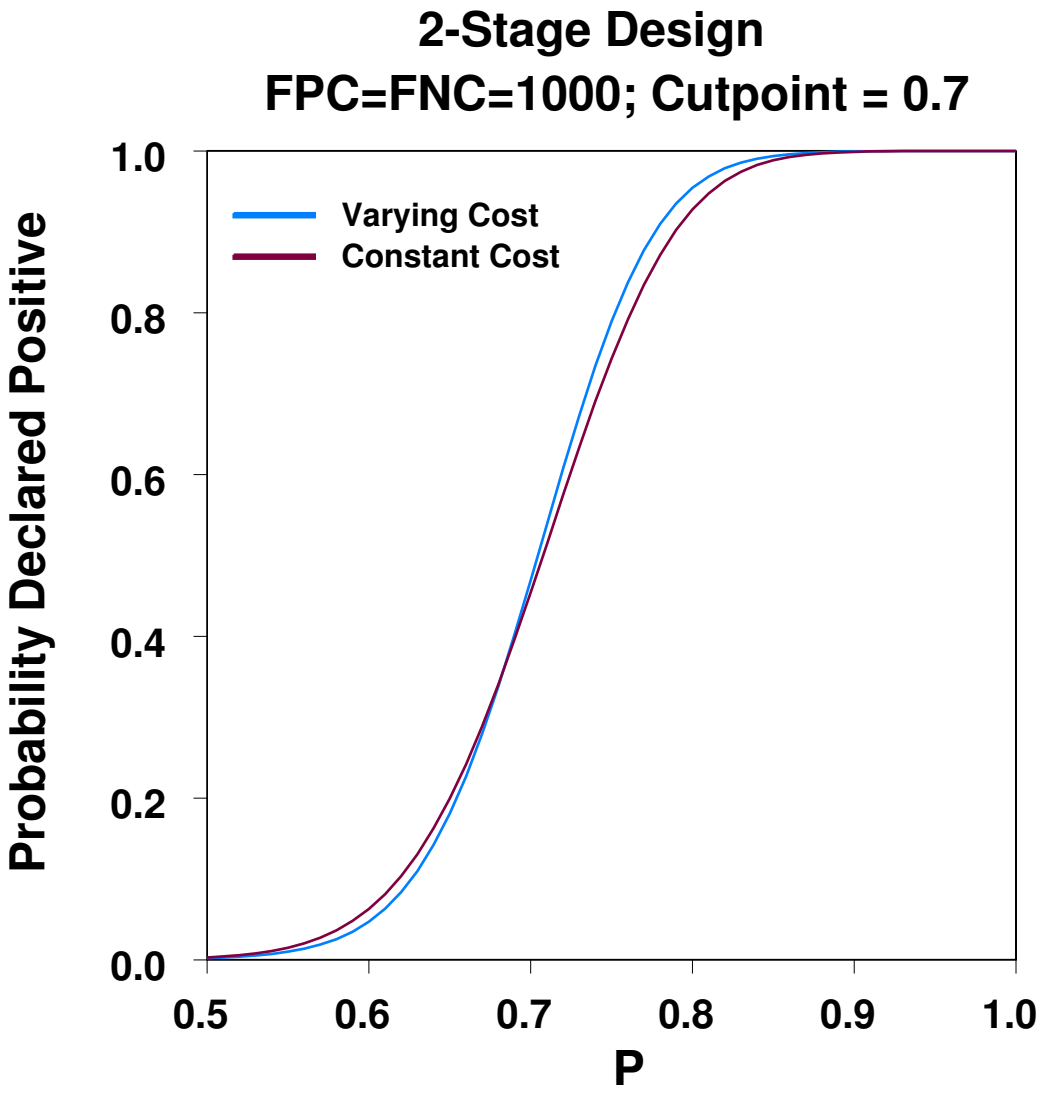
2-Stage Design
FPC=FNC=1000; Cutpoint = 0.7
Probability Declared Positive
1.0
0.8
0.6
0.4
0.2
0.0
Varying Cost
Constant Cost
0.5
0.6
0.7
0.8
0.9
1.0
P

Comments

A unified approach to optimizing true costs, given trial constraints.

- There are a variety of costs and constraints for conducting a trial.
- Decision costs need not merely correspond to type I, type II error probabilities, e.g., distance from cutoff may be significant.
- Our 2-stage design superior to Yao and Venkatraman design, even using their objectives, because our 2nd stage size can depend upon outcome of 1st stage.
- More generally, this approach yields designs with significantly reduced costs because trial structure not artificially restricted.
- 2-stage designs significantly better than 1-stage, and fully sequential better still.
- We provide optimal designs and wide range of exact evaluations.

- Prior researchers noted usefulness of curtailment (early stopping) if responses have variable delays. A program is being developed to optimize designs for such situations.

- In other work, we've incorporated some costs and experimental constraints into clinical trial designs.

- Cost and constraint model may be appropriate for other experimental situations. We are interested in learning about these, as well as general adaptive situations.

References

- J. Hardwick and Q.F. Stout, “Optimal screening designs with flexible cost structures, *Simulation 2001*, (S.M. Ermakov, Yu.N. Kashtanov, and V.B. Melas, eds.), NII Chemistry St. Petersburg, 2001, pp. 253–260. www.eecs.umich.edu/~qstout/pap/Simulation01.pdf
- J. Hardwick and Q.F. Stout, “Using path induction to evaluate sequential allocation procedures” *SIAM J. Scientific Computing* **21** (1999), pp. 67–87. www.eecs.umich.edu/~qstout/pap/SIAMJSC99.pdf
- Y.-G. Wang and D.H.-Y. Leung, “An optimal design for screen trials”, *Biometrics* **54** (1998), pp. 243–250.
- T.-J. Yao, C.B. Begg, and P.O. Livingston, “Optimal sample size for a series of pilot trials of new agents”, *Biometrics* **52** (1996), pp. 992–1001.
- T.-J. Yao and E.S. Venkatraman, “Optimal two-stage design for a series of pilot trials of new agents”, *Biometrics* **54** (1998), pp. 1183–1189.

- You can access all of our papers at
  www.eecs.umich.edu/~qstout/papers.html#adapt