
Adaptive Designs for Balancing Complex Objectives

Janis Hardwick

Quentin F. Stout

University of Michigan

<http://www.eecs.umich.edu/~jphard/>

<http://www.eecs.umich.edu/~qstout/>

Research supported by National Science Foundation

Presented at GlaxoSmithKline April 26, 2002.

OUTLINE

- ✓ Simple problem of determining the better of two therapies



???



- ✓ Discuss potential problem constraints — e.g.

- ★ Costs for time, subject, failures, errors

- ✓ Present response-adaptive approaches:

- ★ using accrued data

- ★ incorporating awareness of future decisions

- ★ optimization vs. competing goals

- ✓ Describe increasingly *intelligent* designs.

- ✓ Results, references and comments → done!

BASIC MODEL

Imagine 2 populations of Bernoulli response data that represent patient responses to treatment arms 1 and 2, (T_1, T_2) .

From T_1 we get $X_{11}, X_{12}, \dots \sim B(1, P_1)$ \searrow

Independent with $(P_1, P_2) \in \Omega = (0, 1) \times (0, 1)$

From T_2 we get $X_{21}, X_{22}, \dots \sim B(1, P_2)$ \nearrow

Assume that we know we want a fixed total sample size of n .

WHAT DO WE MEAN BY DESIGN?

- ★ An experimental “design” is an *algorithm* that specifies how to allocate resources during the study.
- ★ Note that we may require computer *algorithms* to generate the sampling *algorithms* that define our *designs*
- ★ This is the case not only for certain adaptive designs discussed here, but also for such simple considerations as randomization.
- ★ Given a model, an *optimal* design typically addresses a single *criterion*.
- ★ But \longrightarrow it may be desirable to *evaluate* a design on criteria for which it was not *optimized* . . .

We may lose a bit of optimality on 1 criterion to achieve better *trade-offs* overall.

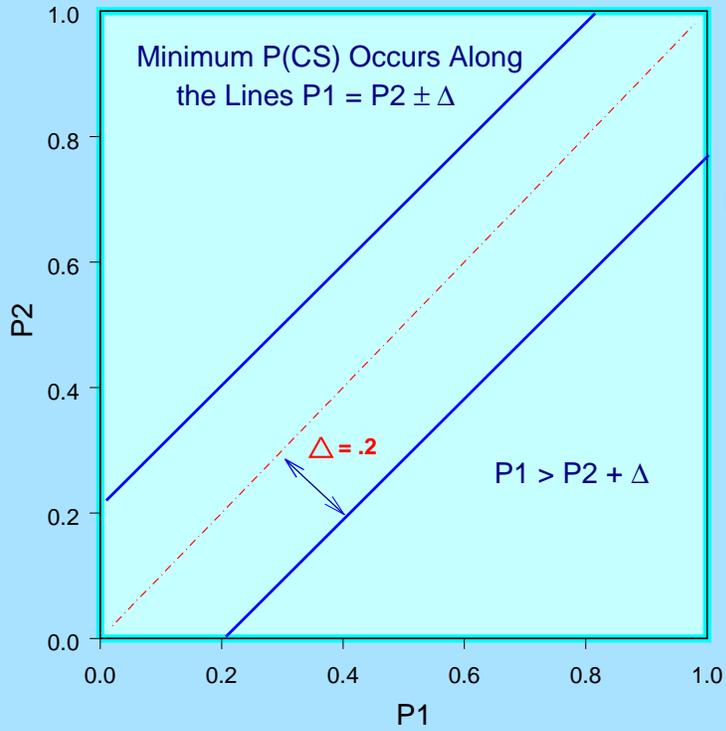
SO WHAT'S ANALYSIS?

- ★ A main part of **analysis** is deciding how to treat the resulting data.
- ★ A key focus of this talk is to show that one can *design* clinical trials using **Bayesian** methods, but still *analyze* the data using **frequentist** perspectives. (*Hardwick and Stout, 1999 "Path Induction"*)
- ★ More generally, **analysis** involves evaluating general operating characteristics of a **design**. (*Error rates, distributions of interesting statistics, robustness to departures from assumptions, etc.*)

Back to our problem . . . SIMPLE GOAL

- ✓ Terminal Decision: Select T_1 or T_2 at end of trial.
- ✓ Goal: Select correctly for $|P_1 - P_2| \geq \Delta$; $\Delta > 0$ a clinically significant difference.
- ✓ In classical allocation designs, it is usual to sample **equally** from the two treatment populations. This is **EQUAL ALLOCATION (EA)**.
- ✓ In general, **no** allocation procedure maximizes power under the alternative $H_A : |P_1 - P_2| \geq \Delta$, for all (P_1, P_2) .
- ✓ One approach is to restrict the notion of an optimal design.
 - ★ $\max \min_{(|P_1 - P_2| \geq \Delta)} \text{P}(\text{Correct Selection})$. *[For this EA is optimal.]*
 - ★ $\mathbf{E}^\xi[\text{P}(\text{Correct Selection})]$, for ξ a prior distribution on (P_1, P_2)

Want High P(Correct Selection) at Terminal Decision



Possible Criteria

MaxMin P(CS: $|P1 - P2| \geq \Delta$)

Integrate P(CS) wrt prior

OR

Pointwise comparisons

Since EA is easy and can be useful for optimizing $P(CS)$ in some special cases \rightarrow

Q: Why extend beyond classical schemes?

A: Most clinical trials have multiple objectives and EA rules may perform arbitrarily badly with respect to other trial criteria:

\rightsquigarrow Incorporate sampling costs: Valuable resources

\rightsquigarrow Minimize length of study: Time

\rightsquigarrow Limit patient suffering during the study: Ethics

\rightsquigarrow Induce balance within groups: Covariates

\rightsquigarrow Reduce variance of estimators: Inference

Alternatives? \longrightarrow *Adaptive Designs*

Q : What are *adaptive* sampling (allocation) designs?

A : Sampling schemes that allow investigators to adjust resource expenditures *while the experiment is being carried out*.

Q : Why is *adaptive* allocation better?

A :  Most interesting objective functions depend on parameters.

 These are unknown or the solution would be trivial.

 Adaptive designs use accruing results to estimate parameters and guide future allocations.

Q: Are there problems associated with *adaptive* designs?

A: \rightsquigarrow Analytically, the data that arise from adaptive designs are less friendly than those we get from fixed sampling designs.

\rightsquigarrow Sample sizes tend to be random variables and the obs are not necessarily independent \implies unknown sampling distributions for standard statistics.

\rightsquigarrow Historically, analysis of adaptive designs has been based on asymptotic approximations \implies may not apply well in practice.

\rightsquigarrow In our work, we generally use computers to carry out exact evaluations of interesting quantities (*Hardwick and Stout, 1995*).

\rightsquigarrow However, simulations are becoming more and more popular as methods for evaluating (*and even designing*) adaptive designs.

A SIMPLE ADAPTIVE DESIGN

- ★ Suppose we're interested in a secondary criterion of *patient survival* during the trial.
- ★ Measure this by $E[\text{Successes Lost}] = [n \times \max\{P_1, P_2\} - \# \text{Failures}]$
- ★ This suggests sampling more often from the better treatment.
- ★ An intuitive way to do this is to **PLAY THE WINNER (PW)**: If the last response was a success, allocate the next patient to the same treatment, otherwise switch.
- ★ Note that **PW** uses only the information in the last observation to make the next decision.

USING MORE INFORMATION

- ★ Another intuitive strategy is called **MYOPIC**: Estimate the unknown parameters using all observed data and allocate the next patient to the treatment that has the highest *expected success rate*.
- ★ **MYOPIC** rules are *aka* One Stage Look Ahead (1SLA) rules.
- ★ Both **MYOPIC** and **PW** strategies have the advantage of being simple to compute.
- ★ But we can do better if we think about the future as well as the past.

TWO-ARM BANDIT PROBLEM

This is basically a problem in optimal learning theory.

- ✓ Slot machine with two arms
- ✓ At pull i , you win $r_i = 0$ or 1
- ✓ Win w.p. P_1 on Arm 1 & P_2 on Arm 2
- ✓ You get n pulls

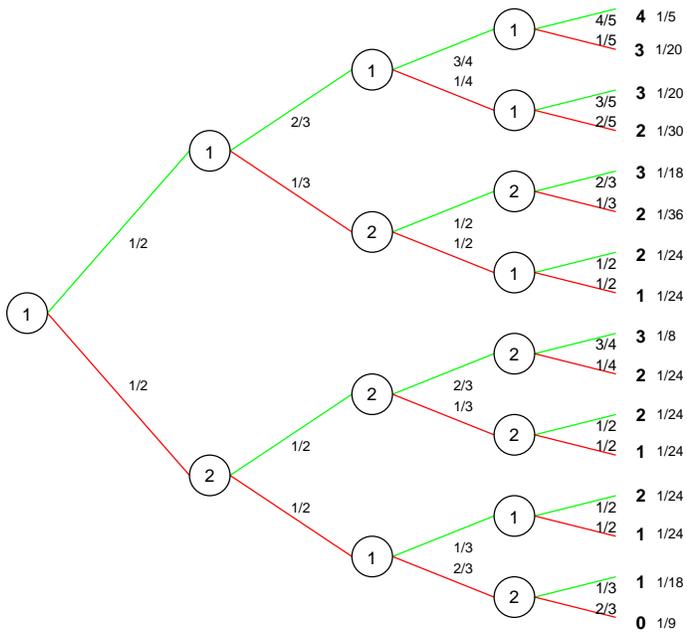
Pull Arm 1 or Arm 2 ??

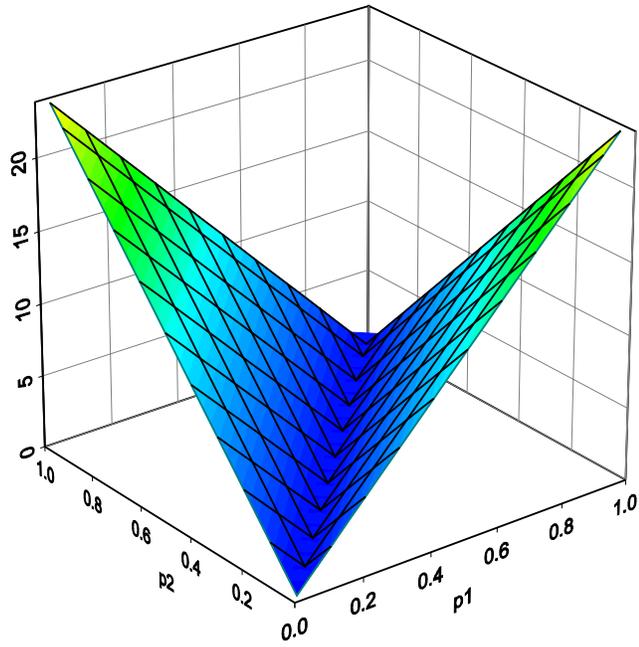
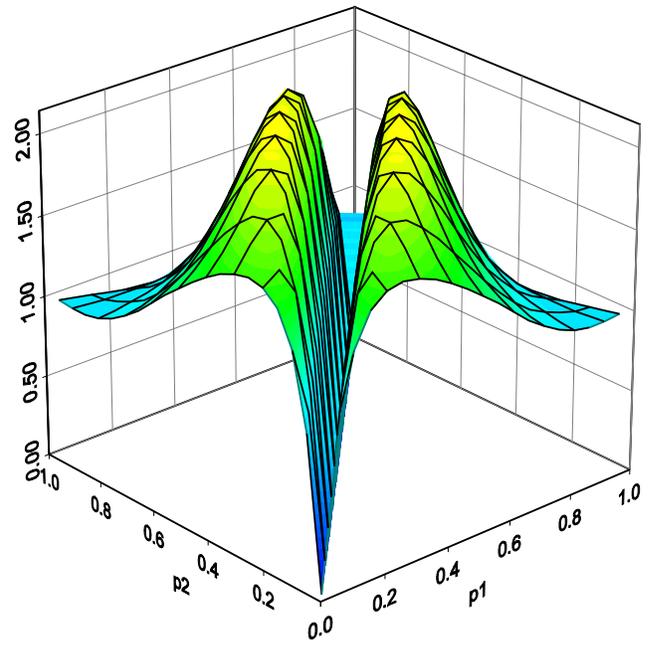


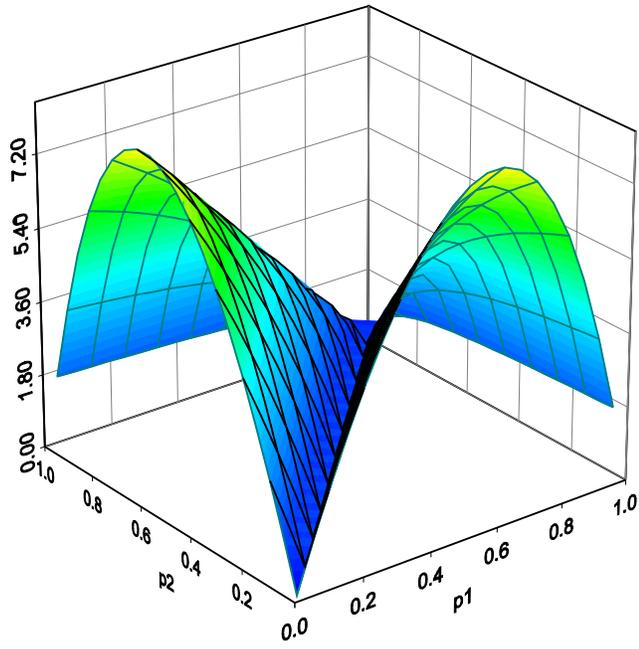
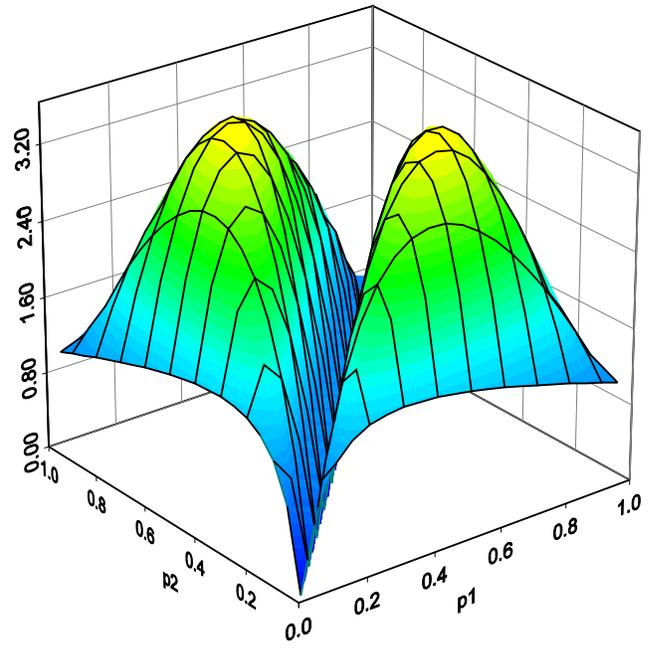
Bandit solution is the optimal solution to the problem of choosing the arms to *maximize* your $\rightarrow E(\text{total reward}) = E(\sum_1^n r_i)$

LOOKING FORWARD

For a sample of size n : bandit solution $\equiv n$ -SLA rule. For $n = 4$:



$E(\text{SUCCESSSES LOST} \mid \text{EA})$  $E(\text{SUCCESSSES LOST} \mid \text{B})$ UNIFORM PRIORS FOR BANDIT; $n=50$

$E[\text{SUCCESSSES LOST} \mid \text{PW}]$  $E[\text{SUCCESSSES LOST} \mid \text{MY}]$  $N=50$

BANDIT TECHNICALITIES

Formulate as Bayesian decision theoretic problem.

Here use beta prior distributions on (P_1, P_2) .

Solution is obtained using Dynamic Programming (DP)

For k arms, DP solution requires computational space and time of $\frac{n^{2k}}{(2k)!}$

SOLVING BANDIT PROBLEMS

Bandit Strategy = **Black Box** ?

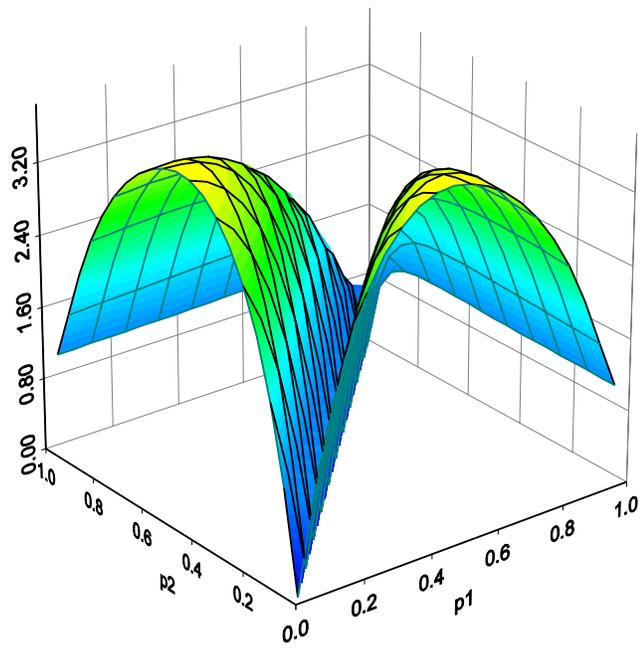
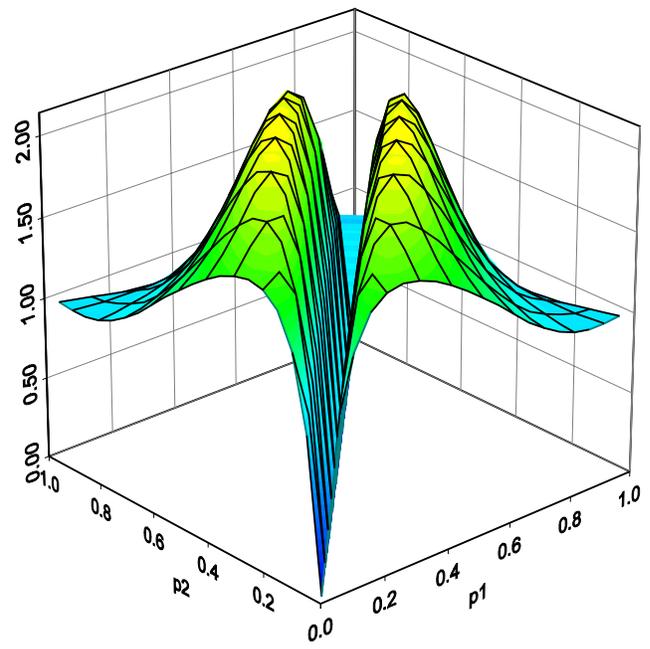
Heuristically \longrightarrow Balancing **immediate gain** (myopic) vs. **information gathering** (exploration for potential future gain)

So, bandits minimize total harm to trial subjects \longrightarrow

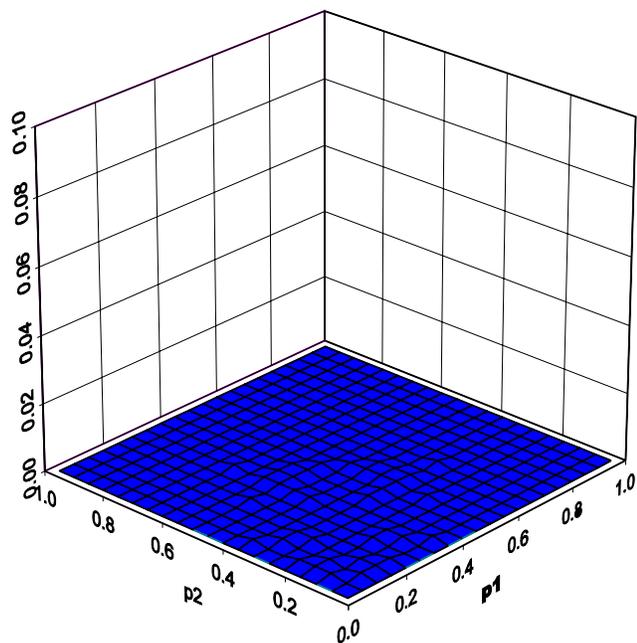
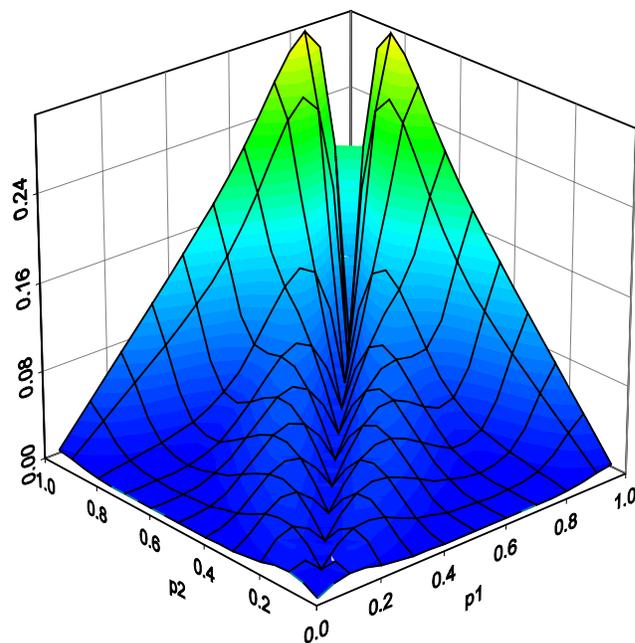
But bandits become **myopic** near the end of a trial, maximizing successes but ignoring desire to also make good decision for post-trial patients.

INCORPORATING FUTURE PATIENTS

- ★ To improve decision making, incorporate reward for **future** patients.
- ★ Assume patient k in the future is “worth” β^k , for some $0 < \beta < 1$
- ★ Elegant theory for maximizing the total “discounted” reward: $E(\sum_1^\infty \beta^i r_i)$
- ★ **Theorem:** For each arm, given prior and observations to date, compute its *Gittins Index* (GI), pull arm of highest index. (*Gittins and Jones, 1974*)
- ★ GI is infeasible to compute. **Lower bound** works well in practice even when applied to finite sample size problems: β acts as control parameter for n .
- ★ Allocation based on the approximate GI is \rightarrow **Modified Bandit** (MB)
(*Hardwick, 1986, 1995*)

$E(\text{SUCCESSES LOST} \mid \text{MB})$  $E(\text{SUCCESSES LOST} \mid \text{B})$ 

UNIFORM PRIORS FOR BANDITS; $n=50$, $\beta = 0.99999999$

$E(\text{PCS} \mid \text{EA}) - E(\text{PCS} \mid \text{MB})$  $E(\text{PCS} \mid \text{EA}) - E(\text{PCS} \mid \text{B})$ 

UNIFORM PRIORS FOR BANDITS; $\beta = .99999999$ FOR MB; $n = 50$

SUMMARY OF DESIGN OBJECTIVES

EQUAL ALLOCATION Maximize information about both arms (maximin), no notion of reward nor of when trial stops

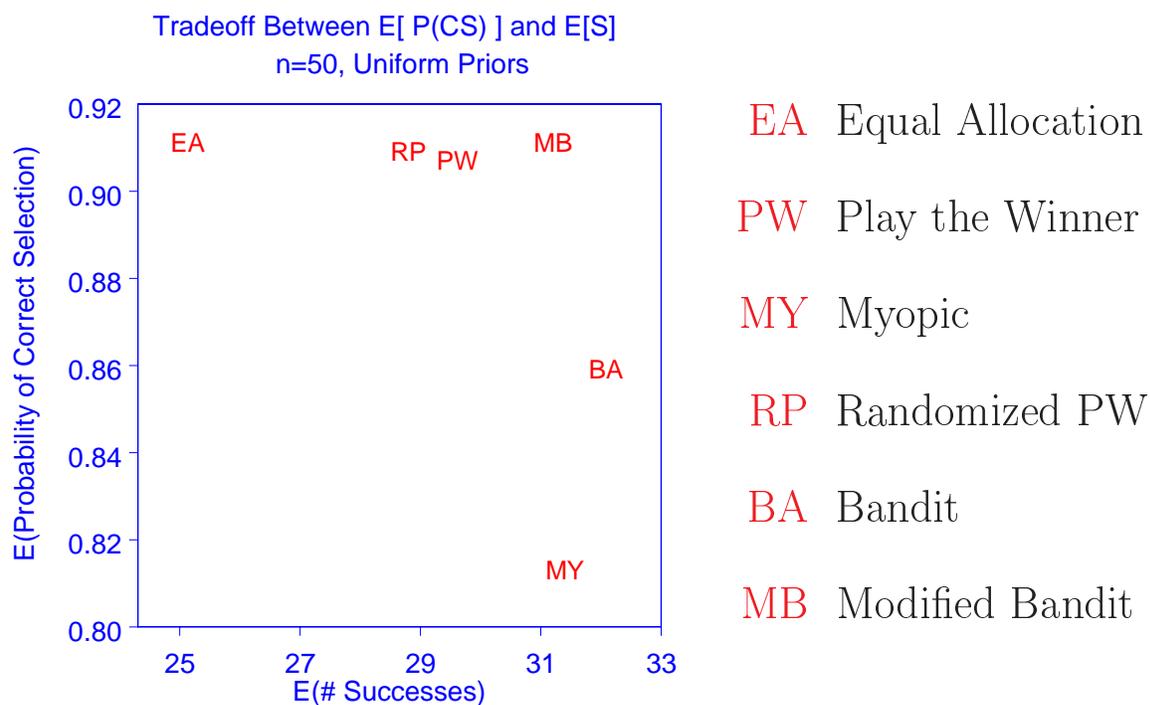
MYOPIC Maximize reward assuming trial stops immediately. No need to gather information.

BANDIT Maximize reward assuming trial is all. Gather enough information to insure greatest total reward in n trials.

MODIFIED BANDIT Maximize reward assuming trial goes on forever, as if to gather information for an infinite future.

SUMMARY MEASURES

Values shown are integrated wrt uniform priors



COMMENTS/FUTURE RESEARCH

Adaptation is good.

- ✓ Does extremely well on more than one objective when not seeking *complete* optimization on one of them.
- ✓ There are numerous ways to approach problems adaptively.
- ✓ Often worth it to trade off optimality for better intuition and design simplicity.
- ✓ Useful to evaluate the designs for robustness.

→ Designs are applicable for many problems. Examples we have worked on include adaptive designs to handle

- ★ Delayed responses
- ★ Censored data
- ★ Sampling in stages
- ★ Dose response models
- ★ Screening models

→ Future research to focus more on developing adaptive designs to handle

- ✓ classical hypothesis testing problems
- ✓ correlated arms — as in dose response models
- ✓ generally more complex models

SOME RELATED LITERATURE

Access our papers via www.eecs.umich.edu/~qstout/papers.html#adapt

- Gittins, J.C. and Jones, D.M. (1974), “A dynamic allocation index for the sequential design of experiments”, *Progress in Statistics* ed. by J. Gani et al., 241–266, North Holland.
- Berry, D.A. and Fristedt, B. (1985), *Bandit Problems: Sequential Allocation of Experiments*, Chapman and Hall.
- Hardwick, J. (1995), “A modified bandit as an approach to ethical allocation in clinical trials”, *Adaptive Designs: Institute Math. Stat. Lecture Notes* **25** 65–87. B. Rosenberger & N. Flournoy, ed.’s.
- Hardwick, J. and Stout, Q.F. (1995), “Exact computational analyses for adaptive designs”, *Adaptive Designs: Institute Math. Stat. Lecture Notes* **25** 223–37. B. Rosenberger & N. Flournoy, ed.’s.

MORE LITERATURE

- Hardwick, J., Oehmke, R. and Stout, Q.F. (1999), “A program for sequential allocation of three Bernoulli populations”, *Comput. Stat. and Data Analysis* **31**: 397–416.
- Hardwick, J. and Stout, Q.F. (1999), “Using path induction for evaluating sequential allocation procedures” *SIAM Journal of Scientific Computing* **21**: 67–87.
- Hardwick, J. and Stout, Q.F. (2000) “Optimizing a unimodal response function for binary variables”, *Optimum Design 2000*, A. Atkinson, B. Bogacka, and A. Zhigljavsky, ed.’s, Kluwer, 2001, pp. 195-208.
- Hardwick, J., Meyer, M.C., and Stout, Q.F. (2001), “Directed walk designs for dose response problems with competing failure modes”, under review.