Discussion of: “Increasing the uptake of long-acting reversible contraceptives among adolescents and young women in Cameroon”

Maximilian Kasy

July 2022

Summary and discussion of the paper

Choosing the objective for adaptive experiments in economics

- Surrogates and value alignment
- Social welfare
- Resource constraints

Summary

- Contraceptive counseling in a hospital in Cameroon.
  - **Treatments**:
    Subsidy rate (1 of 4) for long-acting reversible contraceptives, presentation format (1 of 2).
  - **Primary objective**:
    Minimize probability of unwanted pregnancy for chosen contraceptive, net of the (weighted) cost of subsidy.
- A four stage experimental design:
  1. **Pilot**: Non-adaptive.
     Partition covariate space for targeting, choose tuning parameters.
  2. **Exploration sampling**: Adaptively shift to better-performing treatment, within each of 4 cells, to maximize information for treatment choice.
  3. **Evaluation**: Non-adaptive.
     Most weight on the best policy from stage 2, and the control.
  4. **Follow-up** interviews.

Discussion

- An impressive, pioneering study.
  - Important real-world setting.
  - Balancing numerous objectives:
    Policy choice, participant welfare, several estimands and hypothesis tests.
  - Multi-stage, "human in the loop" procedure.

Q: Motivation for the multi-stage procedure?
  - Could a simpler algorithm achieve the same objectives better?
  - E.g. Exploration sampling throughout, or some similar algorithm?

Q: Why restrict to targeting on 4 sets, based on pilot?
  - No information sharing across cells; coarse representation of context.
  - Alternative: (non-parametric) Bayesian prior $\Rightarrow$ information-sharing across covariate values, and richer targeting?

Discussion

- An impressive, pioneering study.
  - Important real-world setting.
  - Balancing numerous objectives:
    Policy choice, participant welfare, several estimands and hypothesis tests.
  - Multi-stage, “human in the loop” procedure.

Q: Motivation for the multi-stage procedure?
  - Could a simpler algorithm achieve the same objectives better?
  - E.g. Exploration sampling throughout, or some similar algorithm?

Q: Why restrict to targeting on 4 sets, based on pilot?
  - No information sharing across cells; coarse representation of context.
  - Alternative: (non-parametric) Bayesian prior $\Rightarrow$ information-sharing across covariate values, and richer targeting?

Discussion

- An impressive, pioneering study.
  - Important real-world setting.
  - Balancing numerous objectives:
    Policy choice, participant welfare, several estimands and hypothesis tests.
  - Multi-stage, "human in the loop" procedure.

Q: Motivation for the multi-stage procedure?
  - Could a simpler algorithm achieve the same objectives better?
  - E.g. Exploration sampling throughout, or some similar algorithm?

Q: Why restrict to targeting on 4 sets, based on pilot?
  - No information sharing across cells; coarse representation of context.
  - Alternative: (non-parametric) Bayesian prior $\Rightarrow$ information-sharing across covariate values, and richer targeting?

Getting the objective function right

- Key issue here, and for adaptive experiments generally!
  1. Measuring the right outcome.
  2. What about welfare (utility)?
  3. What about resource constraints?
- Remainder of this discussion:
  Elaborating in the context of the present paper, and my own work.

Summary and discussion of the paper

Choosing the objective for adaptive experiments in economics

- Surrogates and value alignment
- Social welfare
- Resource constraints

Measuring the right outcome: Surrogates.

Q: Can we measure what we really care about?

- Unwanted pregnancy rates for the chosen contraceptive $=^?$ actual unwanted pregnancies $=^?$ participant welfare?
- Maximixing short-term formal employment $=^?$ maximizing longer-term employment of any kind?

(As in aur own experiment on ***Job search assistance for refugees in Jordan***.)

A: Under some conditions, yes: "Surrogate outcomes."

Condition: No unobserved causal pathways
to the ultimate outcome of interest.

Athey, S., Chetty, R., Imbens, G. W., and Kang, H. (2019).
The surrogate index: Combining short-term proxies to estimate long-term treatment effects more rapidly and precisely.

Measuring the right outcome: Surrogates.

Q: Can we measure what we really care about?

- Unwanted pregnancy rates for the chosen contraceptive $\stackrel{?}{=}$ actual unwanted pregnancies $\stackrel{?}{=}$ participant welfare?
- Maximizing short-term formal employment $\stackrel{?}{=}$ maximizing longer-term employment of any kind?

(As in aur own experiment on ***Job search assistance for refugees in Jordan***.)

A: Under some conditions, yes: "Surrogate outcomes."

Condition: No unobserved causal pathways
to the ultimate outcome of interest.

Athey, S., Chetty, R., Imbens, G. W., and Kang, H. (2019).
The surrogate index: Combining short-term proxies to estimate long-term treatment effects more rapidly and precisely.

Measuring the right outcome: Surrogates.

Q: Can we measure what we really care about?

- Unwanted pregnancy rates for the chosen contraceptive $\stackrel{?}{=}$ actual unwanted pregnancies $\stackrel{?}{=}$ participant welfare?
- Maximizing short-term formal employment $\stackrel{?}{=}$ maximizing longer-term employment of any kind?

(As in aur own experiment on ***Job search assistance for refugees in Jordan***.)

A: Under some conditions, yes: "Surrogate outcomes."

Condition: No unobserved causal pathways
to the ultimate outcome of interest.

Athey, S., Chetty, R., Imbens, G. W., and Kang, H. (2019).
The surrogate index: Combining short-term proxies to estimate long-term treatment effects more rapidly and precisely.

Measuring the right outcome: Value alignment.

- More generally: Things can go very wrong – "Value alignment."

  Russell, S. (2019).
  Human compatible: Artificial intelligence and the problem of control.

  Thought experiment:

  "The robot programmed to produce as many paperclips as possible, ends up eliminating humanity, since otherwise it could be switched off, which would limit paperclip production..."

- Interesting parallels to contract theory / mechanism design:
  - Picking observable outcome for the algorithm to maximize

  ≈ designing an incentive pay scheme for an agent.

Measuring the right outcome: Value alignment.

- More generally: Things can go very wrong – "Value alignment."

  Russell, S. (2019).
  Human compatible: Artificial intelligence and the problem of control.

  Thought experiment:

  "The robot programmed to produce as many paperclips as possible, ends up eliminating humanity, since otherwise it could be switched off, which would limit paperclip production..."

- Interesting parallels to contract theory / mechanism design:
  - Picking observable outcome for the algorithm to maximize

  $\approx$ designing an incentive pay scheme for an agent.

Measuring the right outcome: Value alignment.

- More generally: Things can go very wrong – "Value alignment."

  Russell, S. (2019).
  Human compatible: Artificial intelligence and the problem of control.

  Thought experiment:

  "The robot programmed to produce as many paperclips as possible, ends up eliminating humanity, since otherwise it could be switched off, which would limit paperclip production..."

- Interesting parallels to contract theory / mechanism design:
  - Picking observable outcome for the algorithm to maximize

  $\approx$ designing an incentive pay scheme for an agent.

What about welfare? Bandits and the social planner.

- Bandit algorithms, exploration sampling:
  Maximize **observable outcomes**.

- Welfare economics, optimal tax theory:
  Maximize **social welfare** – weighted sum of realized utility.
  - Subsidy for contraceptives:
    Why not assume that the Cameroonian women are maximizing their welfare, subject to constraints?

- Tax rate x, demand function G:

$$\text{Social welfare} = \underbrace{x \cdot G(x)}_{\text{Public revenue}} + \lambda \cdot \underbrace{\int_x^1 G(x')dx'}_{\text{Consumer surplus}} .$$

- Welfare at policy x depends on demand for other policies x'!

Adaptive maximization of social welfare

$\Rightarrow$ Exploration needs to take a different form than for bandits.

- Welfare maximizing algorithms:
  - Need to explore more, away from the optimal policy.
  - Worst-case regret rate of $T^{2/3}$, versus $T^{1/2}$ for bandits.
- Work in progress:
  Taking this to the field with the NGO "Mein Grundeinkommen" in Germany:

Adaptive maximization of social welfare

$\Rightarrow$ Exploration needs to take a different form than for bandits.

- Welfare maximizing algorithms:
  - Need to explore more, away from the optimal policy.
  - Worst-case regret rate of $T^{2/3}$, versus $T^{1/2}$ for bandits.
- Work in progress:
  Taking this to the field with the NGO “Mein Grundeinkommen” in Germany:

Adaptive maximization of social welfare

$\Rightarrow$ Exploration needs to take a different form than for bandits.

- Welfare maximizing algorithms:
  - Need to explore more, away from the optimal policy.
  - Worst-case regret rate of $T^{2/3}$, versus $T^{1/2}$ for bandits.
- Work in progress:
  Taking this to the field with the NGO "Mein Grundeinkommen" in Germany:

What about resource constraints?
Combinatorial semi bandits

- Many settings: Resource constraints / matching problems.
- Giving a “treatment” to one unit means we can not give it to another.
  - Monetary budget constraint for contraceptive subsidies?
  - Limited doctor availability?
- Allocations need to be chosen jointly.
- Surprisingly: No cost for worst-case regret rates,
  relative to unconstrained bandits.

Thank you!

References:

Kasy, M. and Sautmann, A. (2021).
Adaptive treatment assignment in experiments for policy choice.

Caria, S., Gordon, G., Kasy, M., Osman, S., Quinn, S., and Teytelboym, A. (2020).
Job search assistance for refugees in Jordan: An adaptive field experiment.

Cesa-Bianchi, N., Colomboni, R., and Kasy, M. (2022).
Adaptive maximization of social welfare.

Kasy, M. and Schwerter, F. (2022).
Adaptive maximization of social welfare: A basic income experiment.

Kasy, M. and Teytelboym, A. (2022).
Matching with semi-bandits.