

Designing basic income experiments

Maximilian Kasy

Department of Economics, Harvard University

April 12, 2019

Introduction


- Suppose one were to run a trial to evaluate a basic income program.
- How should one go about this?

Some questions to answer first:

1. What does “basic income” mean?
2. Why might we want a basic income?
3. What do we expect to learn from basic income trials?
4. And then: How should we design basic income trials?

What does “basic income” mean?

- An unconditional transfer to everyone, regardless of their income?
- A substitute for all other social insurance programs or public goods provision?
- A pathway to the decommodification of our lives and a post-capitalist world?
- **My preferred answer:**
 - A negative income tax,
 - paid upfront, regularly, to individuals,
 - providing a minimum income that no one can fall below,
 - but explicitly taxed away at some rate,
 - and not intended as a substitute to existing programs.



Why would we want a basic income?

- To help us through the coming robot apocalypse, providing sustenance for the superfluous unemployed masses, while a small tech elite runs the world? (“Silicon Valley argument”)
- To replace all public goods provision by cash? (“Chicago argument”)
- To create a post-capitalist utopia where we are liberated from wage labor?
- **My preferred answer:**
 - To create an unconditional safety net, below which no one can fall.
 - To provide outside options, enabling everyone to say “no” to abusive bosses / romantic partners / bureaucrats.
 - To end the intrusive, coercive and expensive surveillance apparatus of current welfare administration.
 - To avoid the repression of wages following from current subsidies of low-wage labor.

What do we expect to learn from basic income trials?

- Whether people who get basic income are
 - happier,
 - healthier,
 - consumer more?

(“Program evaluation approach”)
- Whether basic income
 - discourages work, or
 - encourages investments, or
 - has general equilibrium effects on prices, wages?

(“Empirical public finance approach”)
- **My preferred answer:**
 - To evaluate whether it improves an explicitly specified notion of social welfare, relative to the status quo.
 - To find the specific program parameters that maximize this notion of welfare.

How should we design basic income trials?

- Proof of concept:
 - Give money to a bunch of people.
 - Argue that it was good for them to get the money.
- Conventional program evaluation:
 - Pre-define basic income policy parameters.
 - Split sample equally into treatment and control group, ex ante.
 - Measure a large list of outcomes.
 - Report causal effects of basic income on these outcomes, comparing treatment and control.
- **My preferred answer:**
 1. Embedded in an explicit normative framework, such as the utilitarian welfare framework of optimal tax theory.
 2. Run the experiment in multiple waves, adapting assignment based on the outcomes of previous waves.
 3. Find the policy that maximizes welfare.

Conceptual tools for building an optimal design

- Welfare economics
- Optimal tax theory (Mirrleesian optimal income taxation)
- Machine learning / nonparametric Bayes (Gaussian process priors)
- Adaptive experimental design (Bandits)
- Technometrics (Knowledge gradients)

Kasy, M. (2019). [Optimal taxation and insurance using machine learning – sufficient statistics and beyond.](#)

Journal of Public Economics

Kasy, M. and Sautmann, A. (2019). [Adaptive treatment assignment in experiments for policy choice.](#)

Working Paper

Some references

- **Optimal taxation**

Chetty, R. (2009). [Sufficient statistics for welfare analysis: A bridge between structural and reduced-form methods.](#)

Annual Review of Economics, 1(1):451–488

- **Gaussian process priors**

Williams, C. and Rasmussen, C. (2006). [Gaussian processes for machine learning.](#)
MIT Press

- **Adaptive experiments**

Russo, D. J., Roy, B. V., Kazerouni, A., Osband, I., and Wen, Z. (2018). [A Tutorial on Thompson Sampling.](#)

Foundations and Trends® in Machine Learning, 11(1):1–96

Frazier, P. I. (2018). [A tutorial on Bayesian optimization.](#)

arXiv preprint arXiv:1807.02811

Roadmap

Introduction to optimal taxation

Optimal taxation using machine learning

Experiments for policy choice

Designing basic income experiments

Introduction to optimal taxation

Utility

- General setup:
 - Individual choice set C_i
 - Utility function $u_i(x)$, for $x \in C_i$
 - Realized welfare

$$v_i = \max_{x \in C_i} u_i(x).$$

- Double role of utility
 - Determines choices (individuals choose utility-maximizing x)
 - Normative yardstick (welfare is realized utility)

Can we measure utility?

- Utility can not be observed.
- But we *do* observe choice sets and choices!
- Trick: change the question in two ways
 1. **Changes** in utility, rather than **levels** of utility.
 2. Transfers of **money** that would induce similar changes of utility, rather than changes in **utility** itself.
- \Rightarrow Equivalent variation.

Envelope theorem

- Suppose the prices p_j of various goods change.
- The effect of this change on utility of a given individual i is the same as the effect of a change in her income of

$$dy_i = EV_i = - \sum_j x_{ij} dp_j.$$

- The right hand side is a price index, using the individual's "consumption basket" x_i to weight price changes.
- Put differently: We can ignore behavioral responses to price changes when looking at welfare effects!
- This is the key normative implication of utilitarianism.

Aggregation and disaggregated reporting

- Equivalent variation measures utility changes expressed in monetary units.
- Can aggregate to social welfare, if we have welfare weights:

$$dSWF = \sum_i \omega_i \cdot EV_i$$

- ω_i measures value of an additional \$ for person i
- Could also report welfare changes in a disaggregated way:
 1. Average for various demographic groups, or
 2. average conditional on income.

Redistribution through taxation

- Important policy tool to deal with inequality.
- How to choose a tax and transfer system, tax rates?
- \Rightarrow Theory of optimal taxation.
- Key assumptions:
 1. Evaluate individual welfare in terms of utility.
 2. Take welfare weights as given.
 3. Impose government budget constraint.

Feasible policy changes

- Consider small change in tax rates.
- Has to respect government budget constraint
⇒ Zero effect on revenues.
- Total revenue effect:
 1. Mechanical part: accounting; holding behavior (tax base) fixed.
 2. Behavioral responses: changing tax base.

When are taxes optimal?

- Optimality: no feasible change improves social welfare.
- This implies:
Zero effect on social welfare for any feasible small change.
- \approx First order condition.
- Effect of change on social welfare:
 1. Individual welfare: equivalent variation.
 2. Social welfare: sum up using welfare weights.

Effect on social welfare SWF

- Small change $d\tau$ of some tax parameter.
- Effect on social welfare:

$$dSWF = \sum_i \omega_i \cdot EV_i.$$

- ω_i : value of additional \$ for person i .
- EV_i : equivalent variation.
- By the envelope theorem:
 EV_i is mechanical effect on i 's budget,
holding all choices constant.
- e.g., $EV_i = -x_i \cdot d\tau$ for tax τ on x_i .

Effect on government budget G

- Mechanical effect plus behavioral effect.
- For instance for a tax τ on x_i ,

$$dG = \sum_i x_i \cdot d\tau + dx_i \cdot \tau.$$

- Estimating dx_i part is difficult, the rest is accounting.
- Possible complication: effect of tax change on market prices.
- This complication is often ignored.

Roadmap

Introduction to optimal taxation

Optimal taxation using machine learning

Experiments for policy choice

Designing basic income experiments

Optimal taxation using machine learning

- Standard approach in public finance:
 1. Solve for optimal policy in terms of key behavioral elasticities at the optimum (“sufficient statistics”).
 2. Plug in estimates of these elasticities,
 3. Estimates based on log – log regressions.
- Problems with this approach:
 1. Uncertainty: Optimal policy is nonlinear function of elasticities. Sampling variation therefore induces systematic bias.
 2. Relevant dependent variable is expected tax base, not expected log tax base.
 3. Elasticities are not constant over range of policies.
- Posterior expected welfare based on nonparametric priors addresses these problems.
- Tractable closed form expressions available.

Kasy, M. (2019). [Optimal taxation and insurance using machine learning – sufficient statistics and beyond.](#)

Journal of Public Economics

Optimal insurance and taxation

- Example: Health insurance copay.
- Individuals i , with
 - Y_i health care expenditures,
 - T_i share of health care expenditures covered by the insurance,
 - $1 - T_i$ coinsurance rate,
 - $Y_i \cdot (1 - T_i)$ out-of-pocket expenditures.
- Behavioral response:
 - Individual: $Y_i = g(T_i, \epsilon_i)$.
 - Average expenditures given coinsurance rate: $m(t) = E[g(t, \epsilon_i)]$.
- Policy objective:
 - Weighted average utility, subject to government budget constraint.
 - Relative value of \$ for the sick: λ .
 - Marginal change of $t \rightarrow$ mechanical and behavioral effects.

Social welfare

- Effect of marginal change of t :
 - Mechanical effect on **insurance budget**: $-m(t)$
 - Behavioral effect on **insurance budget**: $-t \cdot m'(t)$
 - Mechanical effect on **utility of the insured**: $\lambda \cdot m(t)$
 - Behavioral effect on **utility of the insured**: 0
By envelope theorem (key assumption: utility maximization)
- Summing components:

$$u'(t) = (\lambda - 1) \cdot m(t) - t \cdot m'(t).$$

- Integrate, normalize $u(0) = 0$ to get social welfare:

$$u(t) = \lambda \int_0^t m(x) dx - t \cdot m(t).$$

Experimental variation, GP prior

- n i.i.d. draws of (Y_i, T_i) , T_i independent of ϵ_i
- Thus

$$E[Y_i | T_i = t] = E[g(t, \epsilon_i) | T_i = t] = E[g(t, \epsilon_i)] = m(t).$$

- Auxiliary assumption: normality, $Y_i | T_i = t \sim N(m(t), \sigma^2)$.
- Gaussian process prior:

$$m(\cdot) \sim GP(\mu(\cdot), C(\cdot, \cdot)).$$

- Read: $E[m(t)] = \mu(t)$ and $\text{Cov}(m(t), m(t')) = C(t, t')$.

Posterior expected welfare

- Denote $\mathbf{Y} = (Y_1, \dots, Y_n)$, $\mathbf{T} = (T_1, \dots, T_n)$, $\mu_i = \mu(T_i)$, $C_{i,j} = C(T_i, T_j)$.
 $\boldsymbol{\mu}$ and \mathbf{C} : vector and matrix collecting these terms.
- Prior moments of welfare:

$$\nu(t) = E[u(t)] = \lambda \int_0^t \mu(x) dx - t \cdot \mu(t), \quad \text{and}$$

$$D(t, t') = \text{Cov}(u(t), m(t')) = \lambda \cdot \int_0^t C(x, t') dx - t \cdot C(t, t').$$

- Notation: $\mathbf{D}(t) = \text{Cov}(u(t), \mathbf{Y} | \mathbf{T}) = (D(t, T_1), \dots, D(t, T_n))$
- Posterior expected welfare:

$$\hat{u}(t) = E[u(t) | \mathbf{Y}, \mathbf{T}] = \nu(t) + \mathbf{D}(t) \cdot [\mathbf{C} + \sigma^2 \mathbf{I}]^{-1} \cdot (\mathbf{Y} - \boldsymbol{\mu}).$$

Application: The RAND health insurance experiment

- Cf. Aron-Dine et al. (2013).
- Between 1974 and 1981, representative sample of 2000 households, in six locations across the US.
- Families randomly assigned to plans with one of six consumer coinsurance rates.
- 95, 50, 25, or 0 percent, 2 more complicated plans (I drop those).
- Additionally: randomized Maximum Dollar Expenditure limits, 5, 10, or 15 percent of family income, up to a maximum of \$750 or \$1,000. (I pool across those.)

Table: Expected spending for different coinsurance rates

	(1)	(2)	(3)	(4)
	Share with any	Spending in \$	Share with any	Spending in \$
Free Care	0.931 (0.006)	2166.1 (78.76)	0.932 (0.006)	2173.9 (72.06)
25% Coinsurance	0.853 (0.013)	1535.9 (130.5)	0.852 (0.012)	1580.1 (115.2)
50% Coinsurance	0.832 (0.018)	1590.7 (273.7)	0.826 (0.016)	1634.1 (279.6)
95% Coinsurance	0.808 (0.011)	1691.6 (95.40)	0.810 (0.009)	1639.2 (88.48)
family x month x site fixed effects	X	X	X	X
covariates			X	X
N	14777	14777	14777	14777


Assumptions

1. **Model:** The optimal insurance model as presented before
2. **Prior:** Gaussian process prior for m , squared exponential in distance, uninformative about level and slope
3. **Relative value** of funds for sick people vs contributors:
 $\lambda = 1.5$
4. Pooling data: across levels of maximum dollar expenditure


Under these assumptions we find:

Optimal copay equals 18%
(But free care is almost as good)


Posterior for m with confidence band



Posterior expected welfare and optimal policy choice



Confidence band for u' and t^*



Roadmap

Introduction to optimal taxation

Optimal taxation using machine learning

Experiments for policy choice

Designing basic income experiments

Experiments for policy choice

The goal of many experiments is to inform policy choices:

1. **Job search assistance** for refugees:
 - Treatments: Information, incentives, counseling, ...
 - Goal: Find a policy that helps as many refugees as possible to find a job.
2. **Clinical trials**:
 - Treatments: Alternative drugs, surgery, ...
 - Goal: Find the treatment that maximize the survival rate of patients.
3. Online **A/B testing**:
 - Treatments: Website layout, design, search filtering, ...
 - Goal: Find the design that maximizes purchases or clicks.
4. Testing **product design**:
 - Treatments: Various alternative designs of a product.
 - Goal: Find the best design in terms of user willingness to pay.

Example

- There are 3 treatments d .
- $d = 1$ is best, $d = 2$ is a close second, $d = 3$ is clearly worse. (But we don't know that beforehand.)
- You can potentially run the experiment in 2 waves.
- You have a fixed number of participants.
- After the experiment, you pick the best performing treatment for large scale implementation.

How should you design this experiment?

1. Conventional approach.
2. Bandit approach.
3. Our approach.

Conventional approach

Split the sample equally between the 3 treatments, to get precise estimates for each treatment.

- After the experiment, it might still be hard to distinguish whether treatment 1 is best, or treatment 2.
- You might wish you had not wasted a third of your observations on treatment 3, which is clearly worse.

The conventional approach is

1. good if your goal is to get a precise estimate for each treatment.
2. not optimal if your goal is to figure out the best treatment.

Bandit approach

Run the experiment in **2 waves**

split the first wave equally between the 3 treatments.

Assign **everyone** in the second (last) wave to the **best performing treatment** from the first wave.

- After the experiment, you have a lot of information on the d that performed best in wave 1, probably $d = 1$ or $d = 2$,
- but much less on the other one of these two.
- It would be better if you had split observations equally between 1 and 2.

The bandit approach is

1. good if your goal is to maximize the outcomes of participants.
2. not optimal if your goal is to pick the best policy.

Our approach

Run the experiment in **2 waves**

split the first wave equally between the 3 treatments.

Split the second wave between

the **two best performing** treatments from the first wave.

- After the experiment you have the maximum amount of information to pick the best policy.

Our approach is

1. good if your goal is to pick the best policy,
2. not optimal if your goal is to estimate the effect of all treatments, or to maximize the outcomes of participants.

Let θ^d denote the average outcome

that would prevail if everybody was assigned to treatment d .

What is the objective of your experiment?

1. Getting precise treatment effect estimators, powerful tests:

$$\text{minimize } \sum_d (\hat{\theta}^d - \theta^d)^2$$

⇒ Standard experimental design recommendations.

2. Maximizing the outcomes of experimental participants:

$$\text{maximize } \sum_i \theta^{D_i}$$

⇒ Multi-armed bandit problems.

3. Picking a welfare maximizing policy after the experiment:

$$\text{maximize } \theta^{d^*},$$

where d^* is chosen after the experiment.

⇒ This talk.

Summary of findings

- **Optimal** adaptive **designs** improve expected welfare.
- Features of optimal treatment assignment:
 - Shift toward better performing treatments over time.
 - But don't shift as much as for Bandit problems:
We have no “exploitation” motive!
- Fully optimal assignment is computationally challenging in large samples.
- We propose a simple **modified Thompson** algorithm (“exploration sampling”).
 - Show that it dominates alternatives in calibrated simulations.
 - Prove theoretically that it is rate-optimal for our problem.

Calibrated simulations

- Simulate data calibrated to estimates of 3 published experiments.
- Set θ equal to observed average outcomes for each stratum and treatment.
- Total sample size same as original.

Ashraf, N., Berry, J., and Shapiro, J. M. (2010). [Can higher prices stimulate product use? Evidence from a field experiment in Zambia.](#)

American Economic Review, 100(5):2383–2413


Bryan, G., Chowdhury, S., and Mobarak, A. M. (2014). [Underinvestment in a profitable technology: The case of seasonal migration in Bangladesh.](#)

Econometrica, 82(5):1671–1748

Cohen, J., Dupas, P., and Schaner, S. (2015). [Price subsidies, diagnostic tests, and targeting of malaria treatment: evidence from a randomized controlled trial.](#)

American Economic Review, 105(2):609–45

Calibrated parameter values



- Ashraf et al. (2010): 6 treatments, evenly spaced.
- Bryan et al. (2014): 2 close good treatments, 2 worse treatments (overlap in picture).
- Cohen et al. (2015): 7 treatments, closer than for first example.

Visual representations

- Compare modified Thompson to non-adaptive assignment.
- Full distribution of regret.
(Difference between $\max_d \theta^d$ and θ^d for the d chosen after the experiment.)
- 2 representations:
 1. Histograms
Share of simulations with any given value of regret.
 2. Quantile functions
(Inverse of) integrated histogram.
- Histogram bar at 0 regret equals share optimal.
- Integrated difference between quantile functions is difference in average regret.
- Uniformly lower quantile function means 1st-order dominated distribution of regret.


Regret and Share Optimal

Table: Ashraf, Berry, and Shapiro (2010)

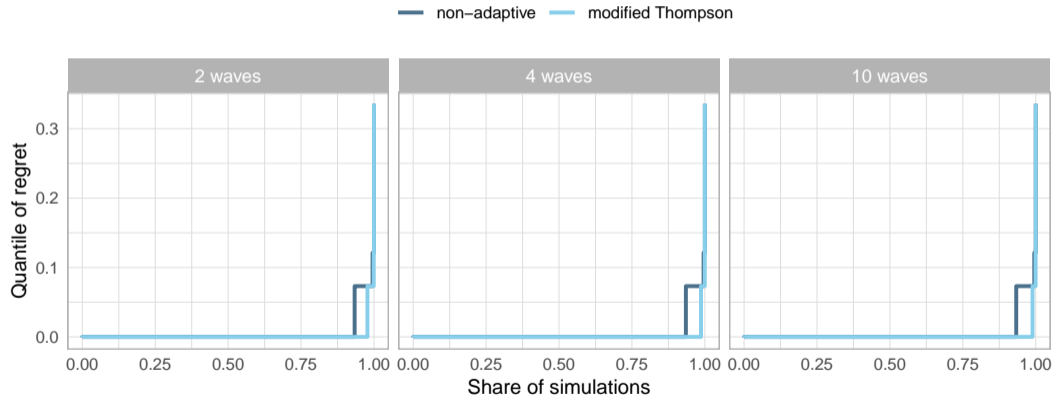
Statistic	2 waves	4 waves	10 waves
Average regret			
exploration sampling	0.0017	0.0010	0.0008
expected Thompson	0.0022	0.0014	0.0013
Thompson	0.0021	0.0014	0.0013
non-adaptive	0.0051	0.0050	0.0051
Share optimal			
exploration sampling	0.978	0.987	0.989
expected Thompson	0.970	0.982	0.982
Thompson	0.972	0.982	0.982
non-adaptive	0.934	0.935	0.933
Units per wave	502	251	100

Policy Choice and Regret Distribution

Ashraf, Berry, and Shapiro (2010)



Policy Choice and Regret Distribution



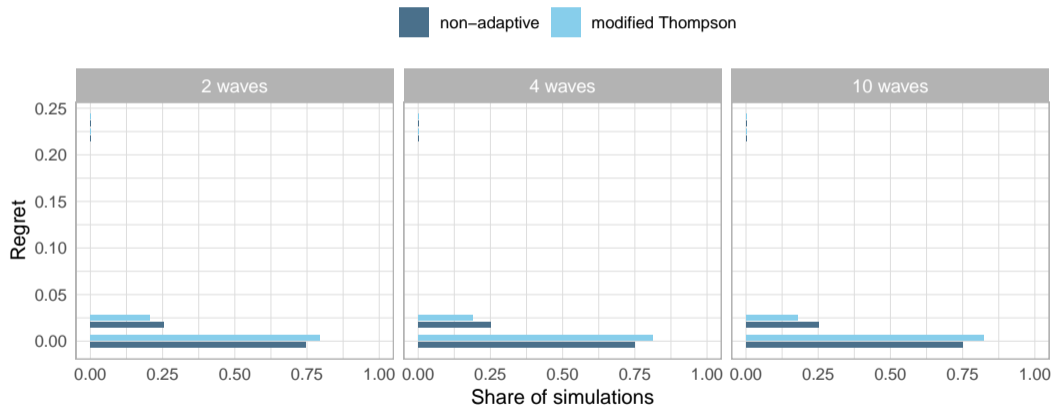
Regret and Share Optimal

Table: Bryan, Chowdhury, and Mobarak (2014)

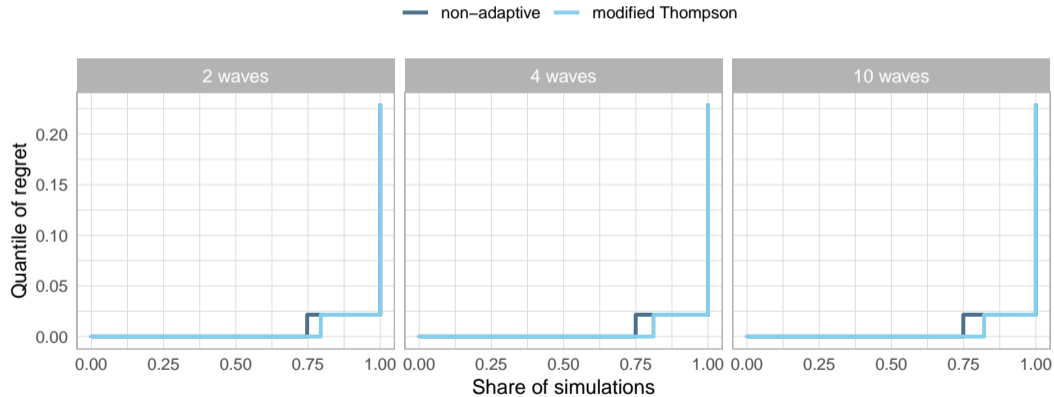
Statistic	2 waves	4 waves	10 waves
Average regret			
exploration sampling	0.0044	0.0041	0.0039
expected Thompson	0.0047	0.0044	0.0043
Thompson	0.0047	0.0044	0.0043
non-adaptive	0.0055	0.0054	0.0054
Share optimal			
exploration sampling	0.794	0.811	0.821
expected Thompson	0.780	0.797	0.800
Thompson	0.781	0.798	0.801
non-adaptive	0.747	0.750	0.749
Units per wave	935	467	187

Policy Choice and Regret Distribution

Bryan, Chowdhury, and Mobarak (2014)



Policy Choice and Regret Distribution




Regret and Share Optimal

Table: Cohen, Dupas, and Schaner (2015)


Statistic	2 waves	4 waves	10 waves
Average regret			
exploration sampling	0.0069	0.0063	0.0060
expected Thompson	0.0074	0.0066	0.0061
Thompson	0.0074	0.0065	0.0062
non-adaptive	0.0087	0.0086	0.0086
Share optimal			
exploration sampling	0.569	0.585	0.592
expected Thompson	0.560	0.579	0.590
Thompson	0.563	0.584	0.590
non-adaptive	0.525	0.526	0.528
Units per wave	1080	540	216

Policy Choice and Regret Distribution

Cohen, Dupas, and Schaner (2015)



Policy Choice and Regret Distribution

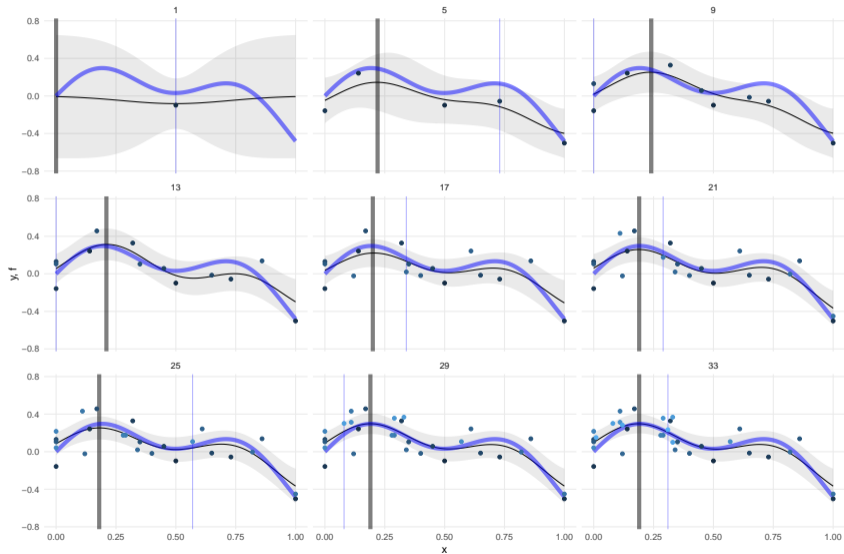


Continuous policy space

The knowledge gradient method

- In basic income experiments, we have a continuous policy space:
Size of basic income, marginal tax rate, ...
- The field of “Bayesian optimization” has developed methods for approximately optimal measurement (treatment assignment) in such settings.
- Knowledge gradient method:
 1. Given outcomes thus far, update the prior for the distribution of the objective function (welfare).
 2. For each possible point of measurement, calculate the prior distribution of the posterior expectation of the objective function.
 3. Assume that after measurement the policy that maximizes expected welfare will be chosen.
 4. Choose the next point of measurement to maximize the expectation of the posterior maximum of welfare.
- \Rightarrow “Greedy knowledge acquisition” targeted at welfare.

The knowledge gradient method - example



Roadmap

Introduction to optimal taxation

Optimal taxation using machine learning

Experiments for policy choice

Designing basic income experiments

Designing basic income experiments

Putting these elements together:

1. Specify welfare weights.
2. Specify the policy space.
(*Variants of basic income.*)
3. Derive mapping from observable outcomes to social welfare.
(*Optimal tax theory.*)
4. Run first wave of experiment.
5. Observe outcomes, update mapping from policies to welfare.
(*Gaussian process priors.*)
6. Pick optimal design points and assignment for second wave.
(*Knowledge gradient.*)
7. Run the next wave, iterate.
8. After the experiment, report the optimal policy, and estimates that allow to calculate the optimal policy for alternative normative choices.

Challenges

- Theoretical:
 1. Generalize mapping from policy parameters to welfare for multi-dimensional policy space.
 2. Set up an appropriate model and non-parametric prior.
 3. Adapt the knowledge gradient method to utilitarian welfare maximization.
- Normative:
 1. Welfare weights: Choosing how much we value marginal \$ for different people.
- Practical:
 1. Measurement: Observing all relevant outcomes, in particular all government transfers received / taxes paid.
 2. Timing: Observing outcomes before assigning next round of treatments.
 3. Complexity: How big should the policy space considered be?
We would like the findings to be easily communicable!

⇒ Exciting work to be done!

Thank you!