

USING ADAPTIVE EXPERIMENTS IN POLICY RESEARCH

Anja Sautmann, DECRG



THE WORLD BANK

IBRD • IDA | WORLD BANK GROUP



**DEVELOPMENT
RESEARCH**

The findings, interpretations, and conclusions expressed in this paper are entirely those of the authors. They do not necessarily represent the views of the World Bank and its affiliate organizations, or those of the Executive Directors of the World Bank or the governments they represent.

Introduction

Randomized control trials (RCTs) represent one of the biggest methodological shifts in public policy research.

- Classical impact evaluations (RCTs, experiments) are designed to estimate the average treatment effect.
- In many policy-making contexts where experiments are used:

~~“What is the effect of a given policy on outcome(s)?”
or “Should this program be implemented?”~~

“Which of these possible alternative programs should we implement for the best outcome?”

“How should this program be implemented?”

→ Goal of analytical work in many operations contexts: among a set of possible approaches, choose those that best achieve project objectives.

Example: what is the best method to contact rice farmers to enroll them into a free agricultural extension service?



Precision Agriculture for Development

- Collaboration with **PAD**
 - Promote custom farming methods in developing countries
 - Here: rice farming in India
- Aim to enroll >1m rice farmers into customized advice service by mobile phone
- Choose between different possible call designs/formats
- Outcome: success = did the farmer pick up the phone and answer enrollment questions?

Logo by PAD.

Example II: How can we design automated calls to inspire parents to practice reading with their first graders?

- Collaboration with NewGlobe's Bridge International Kenya
 - Around 100 schools in Kenya, 2000 first graders
- Goal: get parents to regularly read with their kids at home (2x/week)
- Method: automated calls (IVR) read by a local voice actor.
 - Choose between 6 call variants with different exercise prompts and instructions.
- Outcome: average reading fluency (correct words per minute)

Adaptive Experiments for Policy Research

- **“Which of these possible alternative programs should we implement?”**
 - “Policy Choice Problem”
 - Objective: choose the best out of several program variants *after* the experiment
- Other objectives might be:
 - Achieve best possible outcomes *during* experiment (“multi-arm bandit problem”)
 - Hybrids: balance learning for after the experiment, and giving the best treatment to as many participants as possible
 - Learn treatment effect and best treatment for different types of recipients
... and so on.

Often, a “standard RCT” is not a good fit.

Adaptive Experiments for Policy Research

- **In this talk:** use the “policy choice” example and trial with PAD to demonstrate how adaptive experiments work (Kasy & Sautmann, 2021)
- **Key feature:** carry out the experiment in “waves” – learn as you go – adjust the experiment in later waves to achieve objective as quickly as possible.
- **Growing use of adaptive designs:**
 - Identify and apply successful strategies to help refugees in Jordan find jobs (Caria, Gordon, Kasy, Osman, Quinn, Teytelboym, 2020)
 - *Encourage best-practices adoption for Covid-19 with informational SMS (Bahety et al., 2021)*
 - Increase the uptake of long-acting contraceptives in Cameroon (Athey, Baird, Jamison, McIntosh, Özler, Sama; ongoing)

The Policy Choice Problem

Conducting an experiment in waves to pick one out of several policy (treatment) options.

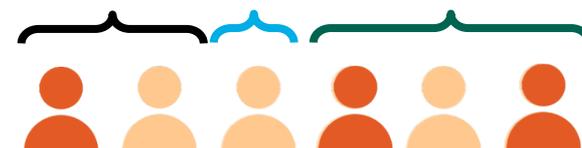
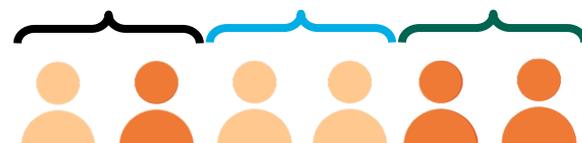
- **Imagine the following situation:**
 - Experiment is (or can be) conducted in waves
 - Three or more treatment options
 - Outcome Y : binary, success (1) or failure (0) [for now]
- *In PAD India example:*
 - *Treatments: contacting mobile phones at different times and with or without a text message sent ahead of time*
 - *Success: the respondent answers a set of enrollment questions.*
- **Policy choice problem: after experiment, we want to pick the option with the highest expected success rate, or equivalently, the highest average Y .**



Precision Agriculture for Development

What does adaptive sampling mean in this setting?

- Assign units to treatment arms at the start of the wave.
- Observe outcomes in each arm.
- Given what was learned, in the next wave, assign units to arms for **optimal continued learning** about what the best option is.
- Observe outcomes again.... etc.



Typically, from wave 2 onwards, it's not the best choice to assign the same number of units to each arm.

Exploration sampling: an algorithm for the optimal(*) adaptive treatment assignment for policy choice.

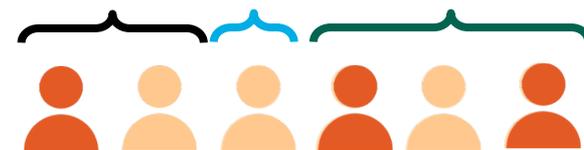
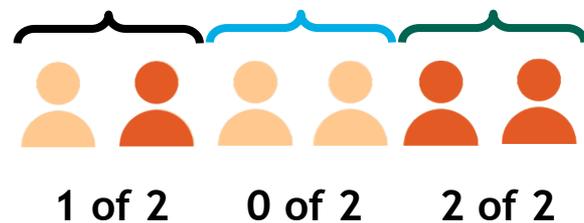
- It is possible to calculate the optimal assignment in each wave exactly.
- However, with larger samples and many treatment arms, it's computationally infeasible to do so.
- **In policy choice problem: use an algorithm called *Exploration Sampling(**)***
 - Approximation of the optimal assignment
 - Builds on current estimated **probability p that a given policy arm is the optimal choice** – assignment proportional to $p(1-p)$.
- ***Exploration sampling* demonstrably improves learning, that is, it picks the best option faster and more reliably.**
 - Different algorithms for other learning objectives.

(*) *Technically, the algorithm is constrained optimal (subject to half the sample being assigned to the best treatment, which in turn is a “robust” condition relative to other choices).*

(**) *Derived from the “multi-arm bandit” literature, which has coined the “exploration-exploitation trade-off.”*

Exploration sampling makes assignment decisions based on Bayesian updating about which treatment has highest success rate.

- Assignment in wave 1 with 3 arms: $p = 1/3$
→ *assign 1/3 at start*
- Observe outcomes
- Calculate probability p for each treatment that it is the optimal option (Bayesian updating) (*)
- Exploration sampling to assign subjects to arms: arm sizes proportional to $p(1-p)$
- Etc.



Key property: since we primarily want to distinguish the best performing options, it's optimal to assign a greater share of units to treatments with high success rates.

* *Note: we are skipping a few details here.*

The Properties of Exploration Sampling

- Bulk of the theoretical paper: optimality properties of exploration sampling as the sample (hypothetically) grows large.
- Simulations show that it performs well in small samples, too.
- As the experiment goes on, the sample shares assigned to different treatment arms converge:
 - Best arm: $\frac{1}{2}$ of the sample
 - Other arms: fixed shares, just right to maximize speed of learning which is the best arm.
- **Results of the form:**
 - “with probability p , arm x has the highest possible outcome” or
 - “the average ‘loss’ from choosing each arm is ... and x has the smallest loss.”



Exploration Sampling in the Real World

Picture: PAD



**DEVELOPMENT
RESEARCH**

What is the best method to contact rice farmers to enroll them into a free agricultural extension service?



Precision Agriculture for Development

- Recall: PAD would like to enroll >1m rice farmers into customized advice service by mobile phone
- Outcome: did the farmer respond to five questions? [Yes = success/1]
- Waves of 600 farmers called through automated service; total of 10K calls
- Six tested treatment arms:
 - Call in the morning or evening
 - For each: 24 hours, 1 hour, or no text message beforehand

Exploration sampling dashboard: results from wave 1 of the experiment



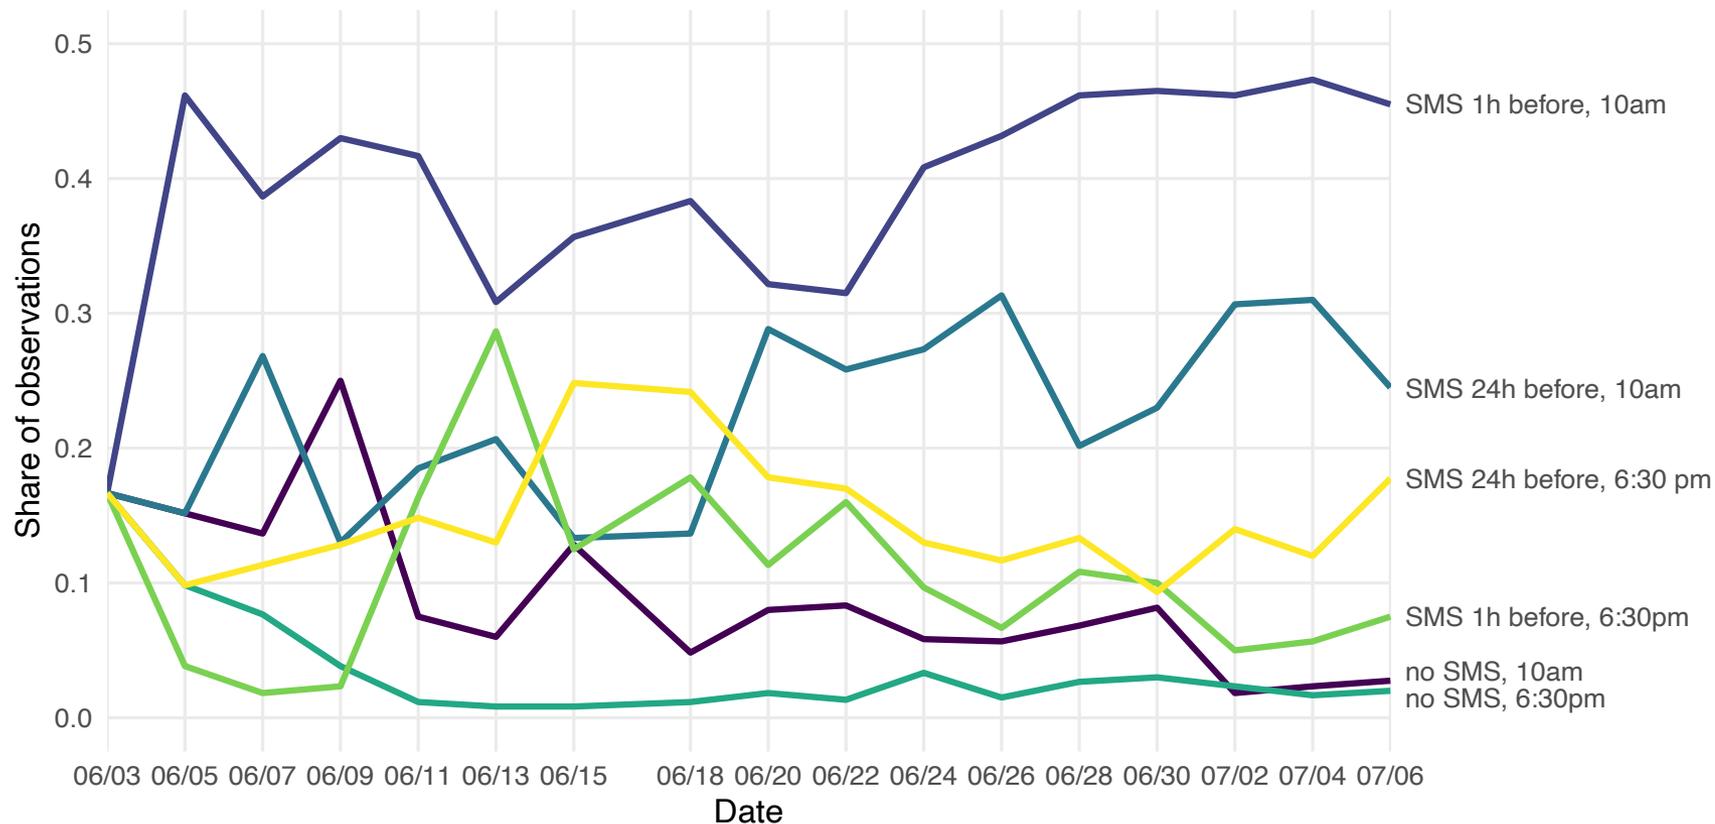
Exploration sampling dashboard: results from wave 2 of the experiment



Exploration sampling dashboard: results at the end of the experiment

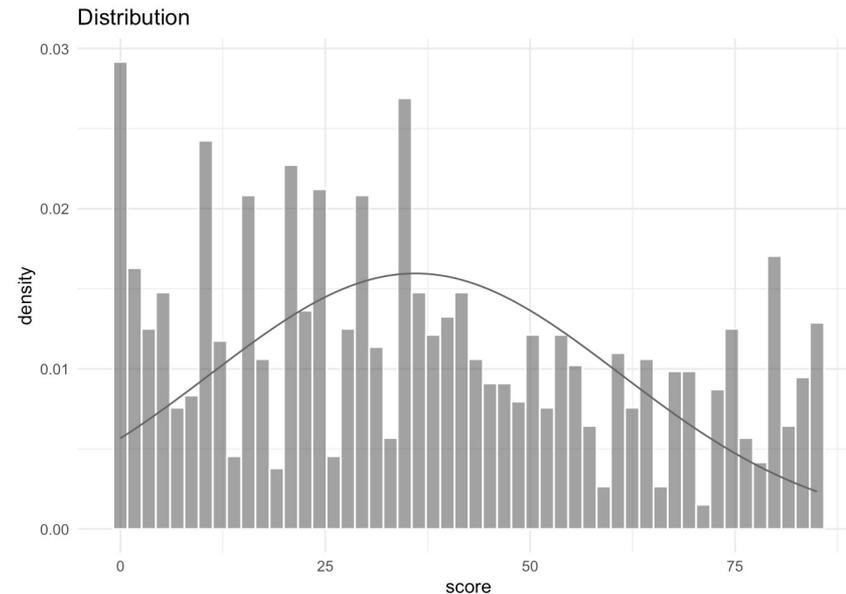


Exploration sampling dashboard: variation in treatment assignment over time



Adaptive experiment in Kenya: promoting home reading practice for improving reading fluency

- First graders in Kenyan schools
- Baseline: endterm exam of term 2
- Biweekly calls with reading exercises from "basic" to "advanced", building on 3rd term text book
- At start of term 3 (May 10): randomly assign half of all children into one of 7 treatment arms, varying:
 - **Difficulty:** "Leveling by baseline" vs. fixed sequence of exercises vs. parental choice
 - **Mode of instruction:** guided by a voice actor on the call, or giving parents exercises to do offline.
 - **Control group** with no calls
- After midterm (June): update "probability optimal" p based on midterm scores; assign second half of the sample using Exploration Sampling.



Conclusion

Adaptive experimental methods can put “tinkering” or piloting on a formal footing.

- Similar to “A/B testing”
- Adaptivity can respond flexibly to past learning

Key advantages:

- **The learning process and resulting decision are replicable and empirically rigorously founded.**
- **For policy choice: more participants benefit from the best treatment options during the experiment.**

What are possible applications?

- **Suitable environments:**
 - outcome realized and observed between waves
 - treatment assignment is flexible from wave to wave
 - For exploration sampling: focus on outcomes *after* the experiment
- **Promising applications:**
 - Education applications: use “wave” structure of grades/terms; particularly important to get to the best treatment in few waves.
 - Technology solutions – treatment assignment pushed centrally to devices; outcomes instantly recorded
 - Intermediate outcomes, such as participation/delivery rate, uptake: quick feedback, uptake is a necessary condition for impact
 - Testing survey methods -- e.g., pilot to find the best question format; then implement the best option in full roll-out

Outlook

- **Ongoing work:** resolve practical open questions, such as how to choose the sample size, using applications like the IVR calls in Kenya
 - Allow for stratification/blocking
 - Deal with unusual outcome distributions, trends, etc.
- An RSB grant to apply exploration sampling, especially in education contexts – with Daniel Rodriguez and Bruno Esposito
- Web-app for trying this out (or implementing it!) at https://maxkasy.shinyapps.io/exploration_sampling_dashboard/

Thank you!

Many thanks to Elizabeth Bond, creator of the representative meeples.

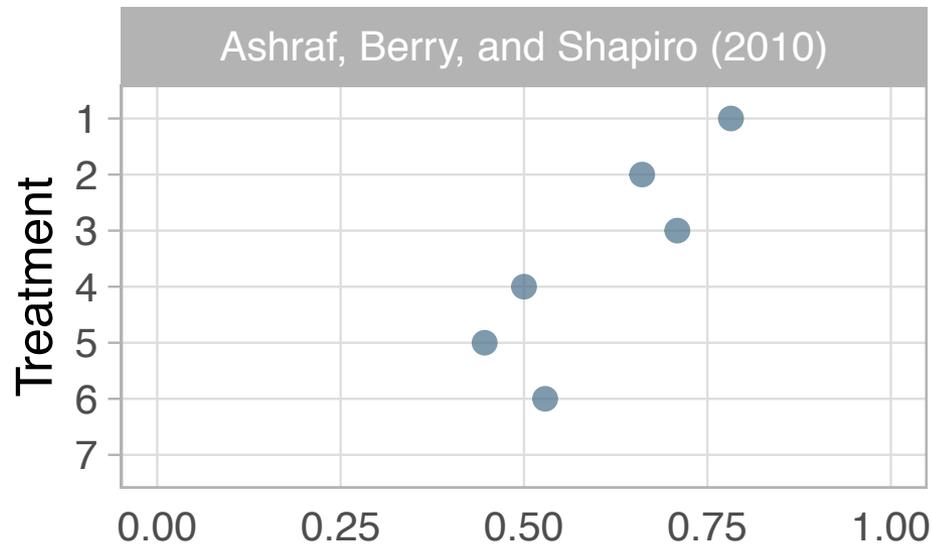


**DEVELOPMENT
RESEARCH**

Appendix: Simulations

Application I: simulations with data from existing experiments to compare performance of different assignment rules.

- Use data from an experiment with several treatments and pretend it is the “truth” in order to test exploration sampling.
- Ashraf, Berry, and Shapiro (2010):
 - Vary price of water disinfectant Clorin from Kw 300 to Kw 800
 - Outcome of interest: probability of purchase



Exploration sampling performs best among assignment methods, and performance increases as the number of waves increases.

Ashraf, Berry, and Shapiro (2010)

Statistic	2 waves	4 waves	10 waves
Average policy regret			
exploration sampling	0.0017	0.0010	0.0008
expected Thompson	0.0022	0.0014	0.0013
non-adaptive	0.0051	0.0050	0.0051
Units per wave	502	251	100