

# Preference Learning for Stakeholder Management

Peter Frazier

Cornell & Uber



# How testing helped Cornell University become a model of COVID-19 prevention

2 years ago | [News](#) | 8:19 

At the start of the school year, Cornell University implemented a strategy of regular testing and robust contact tracing on campus. The plan was expensive, but it's prevented any major COVID-19 outbreaks at the New York institution.



Cornell COVID-19 Mathematical Modeling Team, Summer 2020:



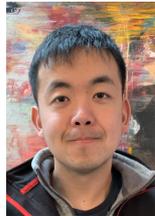
J. Massey  
Cashore



Ning  
Duan



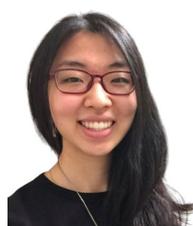
Alyf  
Janmohamed



Brian  
Liu



Jiayue  
Wan



Yujia  
Zhang



Shane  
Henderson



David  
Shmoys

OPINION | COMMENTARY

## Why Cornell Will Reopen in the Fall

Students will return to Ithaca in any case. On campus, we can track and isolate Covid cases.

By Michael I. Kotlikoff and Martha E. Pollack  
June 30, 2020 12:45 pm ET

🔖 SAVE 🖨️ PRINT 🗨️ TEXT

229 



Cornell University in Ithaca, N.Y.

PHOTO: UNIVERSAL IMAGES GROUP VIA GETTY IMAGES

Consider two scenarios. University A decides to reopen. For the health and safety of students, faculty and staff, it institutes a screening program to identify asymptomatic students infected with the novel coronavirus and prevent them from spreading it by repeated testing and isolation. The school also monitors symptoms daily, restricts group sizes, modifies classrooms and dorms, secures extensive quarantine capacity, restricts travel, and imposes requirements for masks and social distancing.

University B decides that this is too risky and chooses to play it safe. The school doesn't reopen for residential instruction this fall and opts instead to teach all courses online. It takes cautious steps to open for selected professional or graduate programs and research efforts, but doesn't implement the complex process of screening thousands of undergraduates and modifying the learning environment for social distancing.

Surprisingly, [epidemiological modeling](#) done by a group led by Cornell Prof. Peter Frazier suggests that despite playing it safe, sometime during the fall University B may well experience markedly worse health outcomes in its community, while University A will have more effectively safeguarded public health.

NEWSLETTER SIGN-UP

### Opinion: Morning Editorial Report

All the day's Opinion headlines.

PREVIEW

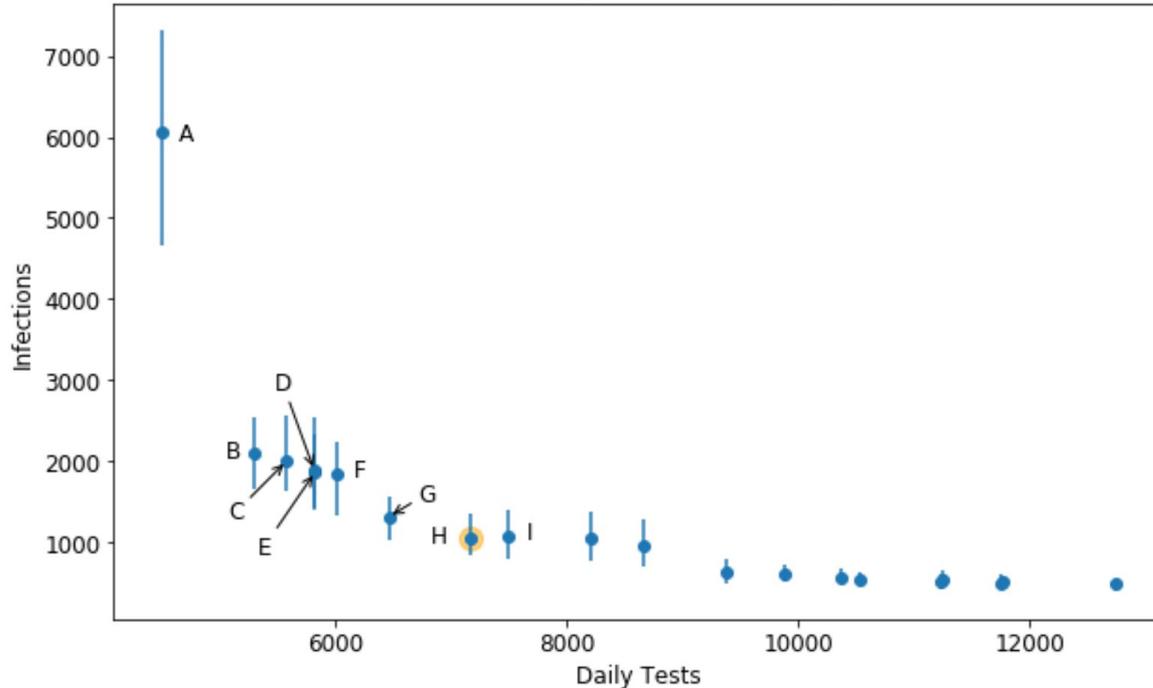


SUBSCRIBE

For many universities, closing the campus to undergraduates is probably not the safest option—notwithstanding concerns that college

# We used this graph to help the provost & lab choose testing frequencies

Infections vs. Daily tests for various policies



- Each letter corresponds to a collection of test frequency assignments
- E.g., “B” means:  
UG in high-density housing 2x / week  
Off-campus staff 1x / month  
Everyone else 1x / week
- Y-axis (infections) are predicted by a simulation-based from the Cornell COVID-19 mathematical modeling team
- Yellow highlight shows the one chosen

Successful OR practice requires  
**understanding stakeholders**

# We took these steps to help the provost & lab choose testing frequencies

- Listen to stakeholders to understand their goals & beliefs
- Run analysis to identify a set of options likely to satisfy them
- Present stakeholders with this set and let them choose

# Can AI help with these things?

- Listen to stakeholders to understand their goals & beliefs
- Run analysis to identify a set of options likely to satisfy them
- Present stakeholders with this set and let them choose



Astudillo & F., "Multi-attribute Bayesian Optimization with Interactive Preference Learning", AISTATS 2020

Lin, Astudillo, F., Bakshy, "Preference Exploration for Efficient Bayesian Optimization with Multiple Outcomes" AISTATS 2022

# Bayesian Optimization is a black-box derivative-free non-convex optimization method

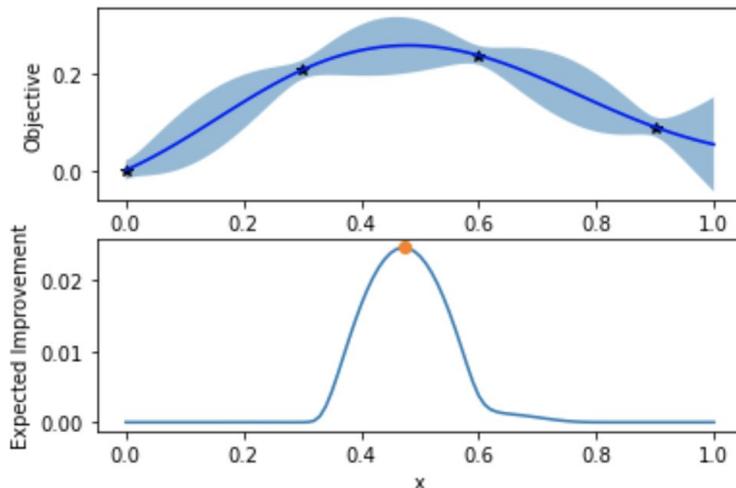
- $x$  = decision to be made
- Slow computer code or experiment can predict objective function  $f(x)$
- $f(x)$  has been evaluated at some  $x$
- Bayesian ML predicts at other  $x$

**Acquisition Function:** the value of a slow evaluation, e.g., the “expected improvement”

$$EI(x) = E[\max\{f(x) - f^*, 0\}]$$

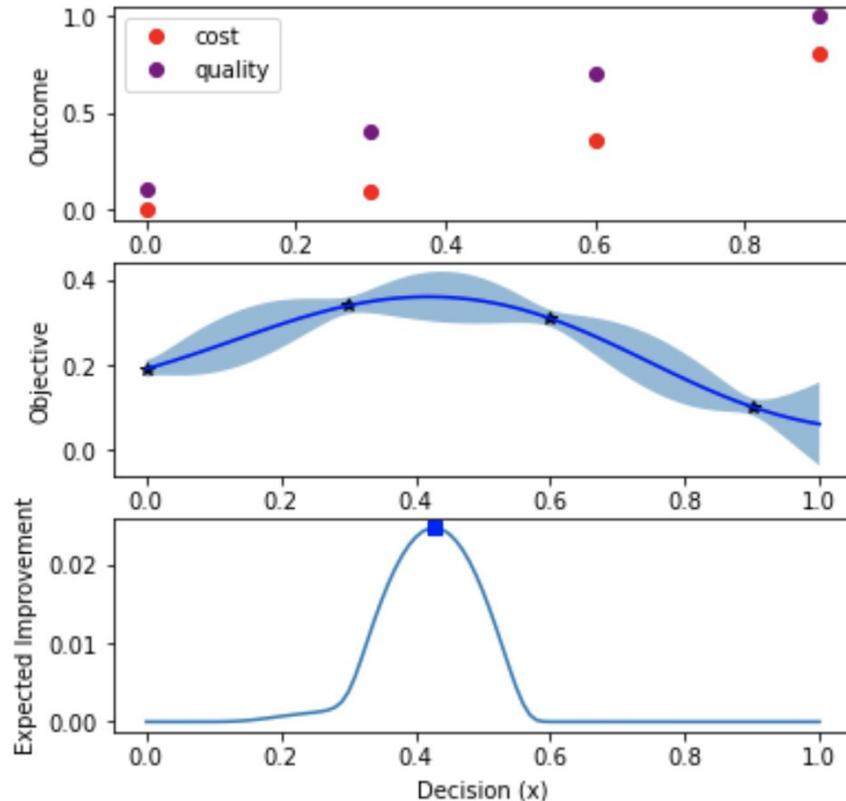
where  $f^*$  is the best  $f(x)$  seen so far.

**Bayesian optimization** runs a slow experiment at the  $x$  with the largest acquisition function value, refits the ML model, & repeats



# Sometimes the objective is the utility of an outcome vector

- Slow experiment gives outcome vectors  $[h_1(x), \dots, h_k(x)]$
- $u$  = utility function over outcomes
- Objective is  $f(x) := u(\mathbf{h}(x))$
- Example:
  - $\mathbf{h}(x) = [\text{quality}(x), \text{cost}(x)]$
  - $u(\mathbf{h}(x)) = \text{quality}(x) - \text{cost}(x)$

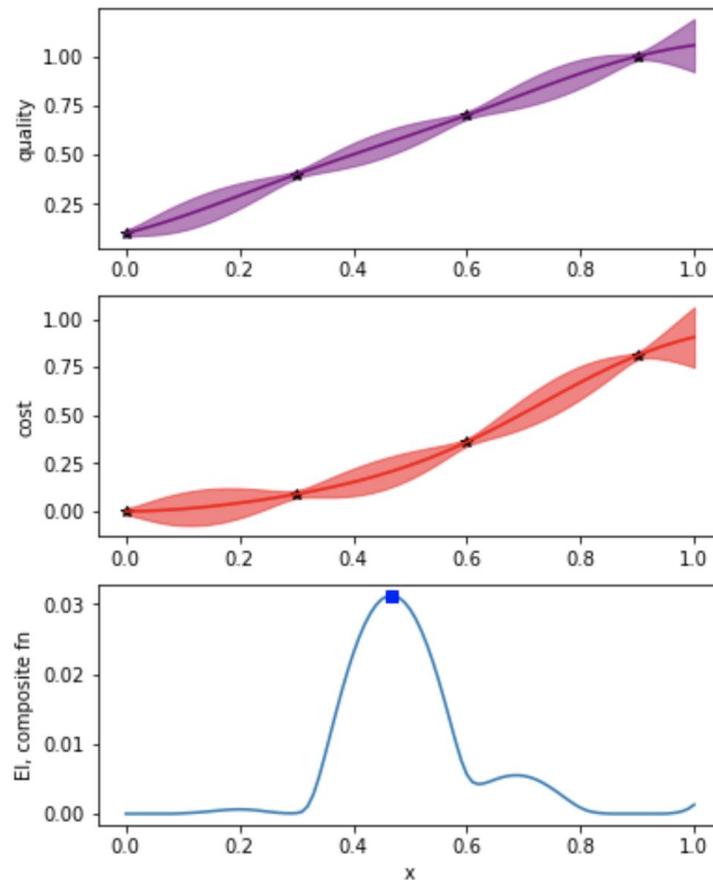


# Sometimes the objective is the utility of an outcome vector

You can also put a Bayesian ML model on  $\mathbf{h}(x)$  and use  
 $EI(x) = E[\max\{u(\mathbf{h}(x)) - f^*, 0\}]$

Example:

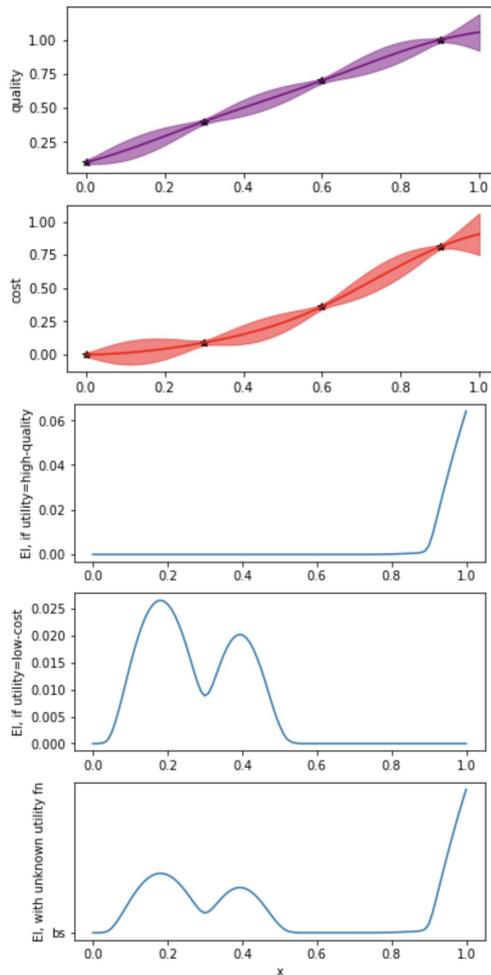
- $\mathbf{h}(x) = [\text{quality}(x), \text{cost}(x)]$
- $u(\mathbf{h}(x)) = \text{quality}(x) - \text{cost}(x)$



Astudillo & F. "Bayesian optimization of composite functions" ICML 2019,

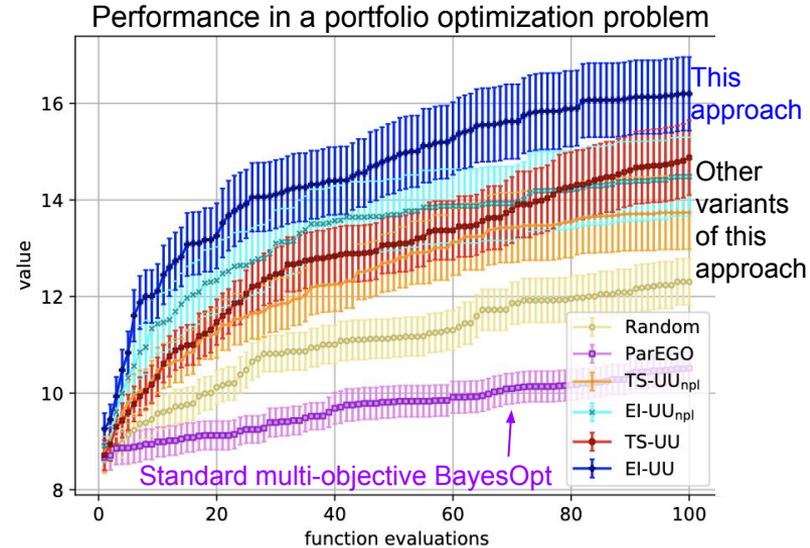
# If the utility function is unknown, we learn it with Bayesian preference learning

- Suppose we don't know the utility function
- Preference learning (for learning the utility fn):
  - Present the stakeholder pairs of outcome vectors. Ask, "Which do you prefer?"
  - Learn a Bayesian ML model on the utility function
  - Active learning decides which pairs to present
- Bayesian optimization (for choosing  $x$  to eval):
  - Prediction: Fit a GP to outcomes  $\mathbf{h}(x)$
  - Acquisition function:  $EI(x) = E[\max\{u(\mathbf{h}(x)) - f^*, 0\}]$  where the expectation  $E$  is over the random utility function  $u(\cdot)$  and the outcome vector  $\mathbf{h}(x)$
  - Run the slow computer code where  $EI(x)$  is largest
- Iterate between asking the stakeholder questions & running the slow computer code



# This approach provides lots of value compared to multi-objective optimization

- In use at Meta for product improvement:
  - Stakeholder = product manager
  - Slow experiment = A/B experiment
  - Utility = Quality of Instagram, Facebook, etc.
- Can handle many outcomes; multi-objective optimization usually limited to 3 outcomes
- Can leverage preferences learned from active queries
- Can leverage passive stakeholder observations, e.g., choices in related decision problems



# There are many kinds of stakeholder engagement

## Dean of Faculty

Office of the Dean

2020-21 August 5

[Home](#) > [Faculty Senate](#) > [Actions and Archives](#) > [Archived Senate Meetings](#) > [2020-21 August 5](#) > [Updates from the Modeling Team](#)

## Updates from the Modeling Team

INDEPENDENT SINCE 1880

# The Cornell Daily Sun

COVID-19 Briefing News Opinion Sports Arts & Culture Science Dining Multimedia Specials Sunspots 4/20



In their plan for reopening, President Pollock and Provost Kotlikoff cite modeling that suggests a hybrid semester, with stringent testing, is safer than an entirely remote semester. The Sun interviewed the professor leading that modeling effort, Prof. Peter Frazer.

CAMPUS REOPENING July 2, 2020

### What Happens if Someone Throws a Party? Questions and Answers on the Model Reopening Campus

By Ariel Oza



## ANNA KELLES

for NY State Assembly District 125

POLICY FOR THE PEOPLE

Kelles4NY.com



**GMA COVER STORY**  
**COVID-19**  
**CORONAVIRUS ON CAMPUS**  
EXCLUSIVE LOOK AT CORNELL UNIVERSITY'S MASSIVE TESTING EFFORT

abc

# Lots of room for innovation in AI-enabled stakeholder engagement for OR applications

- Understand stakeholder goals, beliefs & incentives
- Understand how groups of stakeholders influence each other
- Predict how stakeholders will react to communication
- Managing trust (in the OR analyst and her models)