Efficient and Targeted COVID-19 Testing via Reinforcement Learning

Kimon Drakopoulos

Simons Institute, November 2022

Joint work with **Hamsa Bastani, Vishal Gupta, Jon Vlachogiannis,**
Christos Hadjicristodoulou, Pagona Lagiou, Gkikas Magiorkinis, Dimitrios Paraskevis, Sotirios Tsiodras

4 questions about testing

1. Are **perfect** tests worth waiting for?

2. Is **higher accuracy** better for social **welfare**?

tests as products

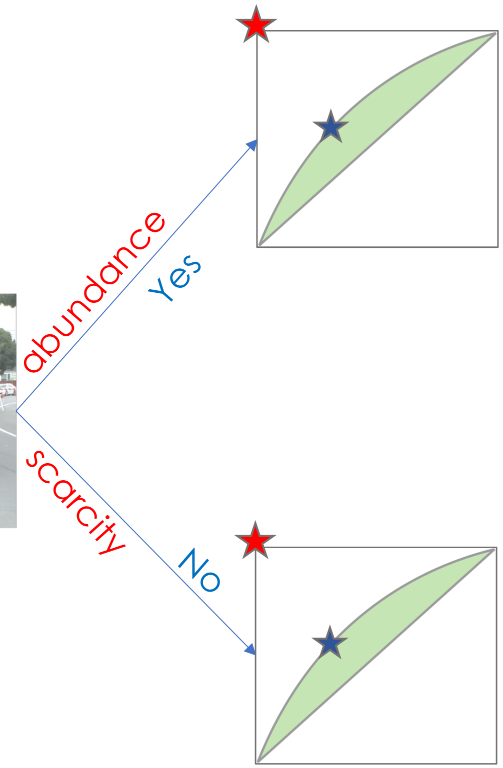

Why Perfect Tests May Not be Worth Waiting For: Information as a Commodity
(K.D., R. Randhawa, *Published at* ***Management Science*** *FastTrack*)

4 questions about testing

3. How to use **Data** and **Operations** to mitigate **pandemics:**
   - **a case study on border control**

4. Can **public** data be used for effective **mitigation** at the border?
   - **not for defining travel protocols**

Implementation and Supply chain

- 300 doctors and nurses
- 200 policemen and firemen
- 32 labs
- ~ millions visitors

COVID-19-EVA Executive Committee (Greece)

- Kyriakos Mitsotakis, Prime Minister
- Nikos Hardalias, Deputy Minister of Civil Protection
- Panagiotis Prezerakos, Secretary General of Public Health

Nature (2021)

Problem in hand

100 to **120 MILLION**
direct tourism **jobs at risk**
(UNWTO)

- Greece **conservative** approach:
  - lockdown of different stringency levels March-May, 2020

- Tourism **25% GDP**

- Economy heavily affected by lockdown
  - Small IT sector
  - Small manufacturing sector

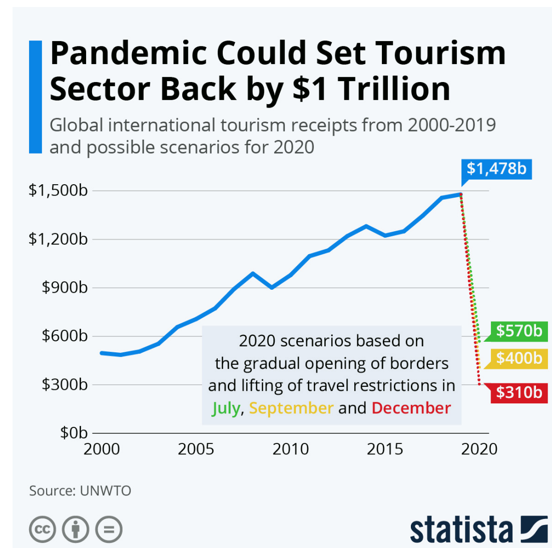

Not really choice on if, choice on how

(6) On 11 June 2020, the Commission adopted a Communication[6] which recommended to extend the restriction on non-essential travel into the EU until 30 June 2020, and which sets out an approach for a gradual lifting of the restriction on non-essential travel into the EU as of 1 July 2020. All Member States have implemented the further extension until 30 June.

Can be done effectively in a data driven way

(10) Decisions on the possible lifting of the restriction on non-essential travel into the EU should take into account the epidemiological situation within the EU, i.e. the average number of COVID-19 cases over the last 14 days and per 100 000 inhabitants.

Not effective advice

EU Council Recommendation on the temporary restriction on non-essential travel into the EU and the possible lifting of such restriction

Pre-Summer 2020: Travel back in time…

- **<u>No vaccines available</u>**

- Testing was scarce
  - Rapid tests unreliable
    - Especially for asymptomatic infections
  - PCR testing requires specialized machinery
    - 3-6 month lead time for new machines

- knowledge was limited
  - How long after exposure until contagious?
  - How long asymptomatic and contagious?

Before we joined (before May 1st)

- **Random testing** at the border (1 in 10-15-open loop)
- **Grey-listing** (72-hour PCR test required) "as we go" based on 14-day notification rate.
- **Red-listing** practically unimplementable with arrivals ~ millions and GDPR constraints.

Eva: System Overview

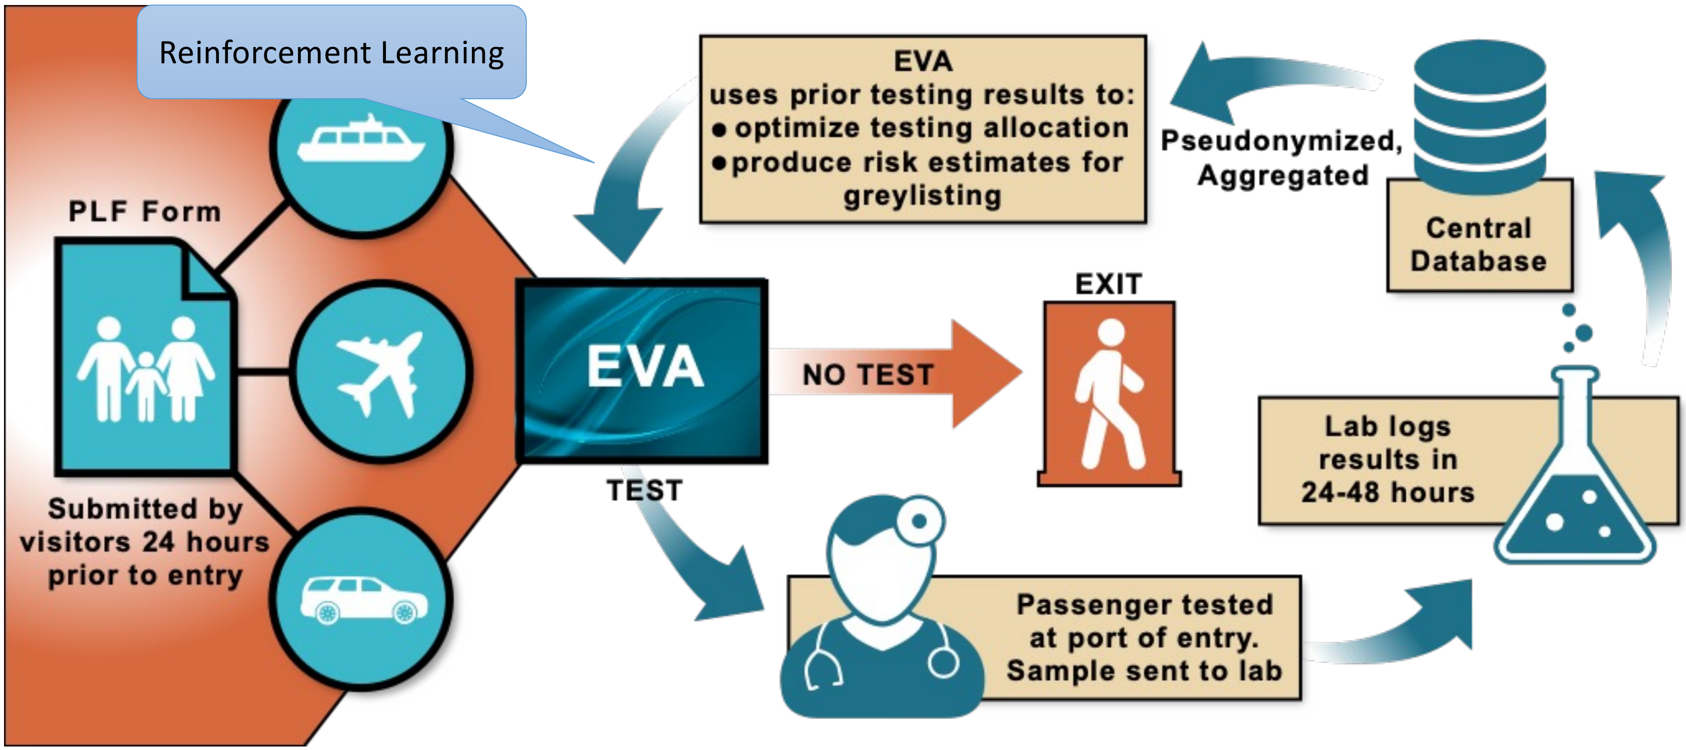

Eva is more than a reinforcement learning algorithm!

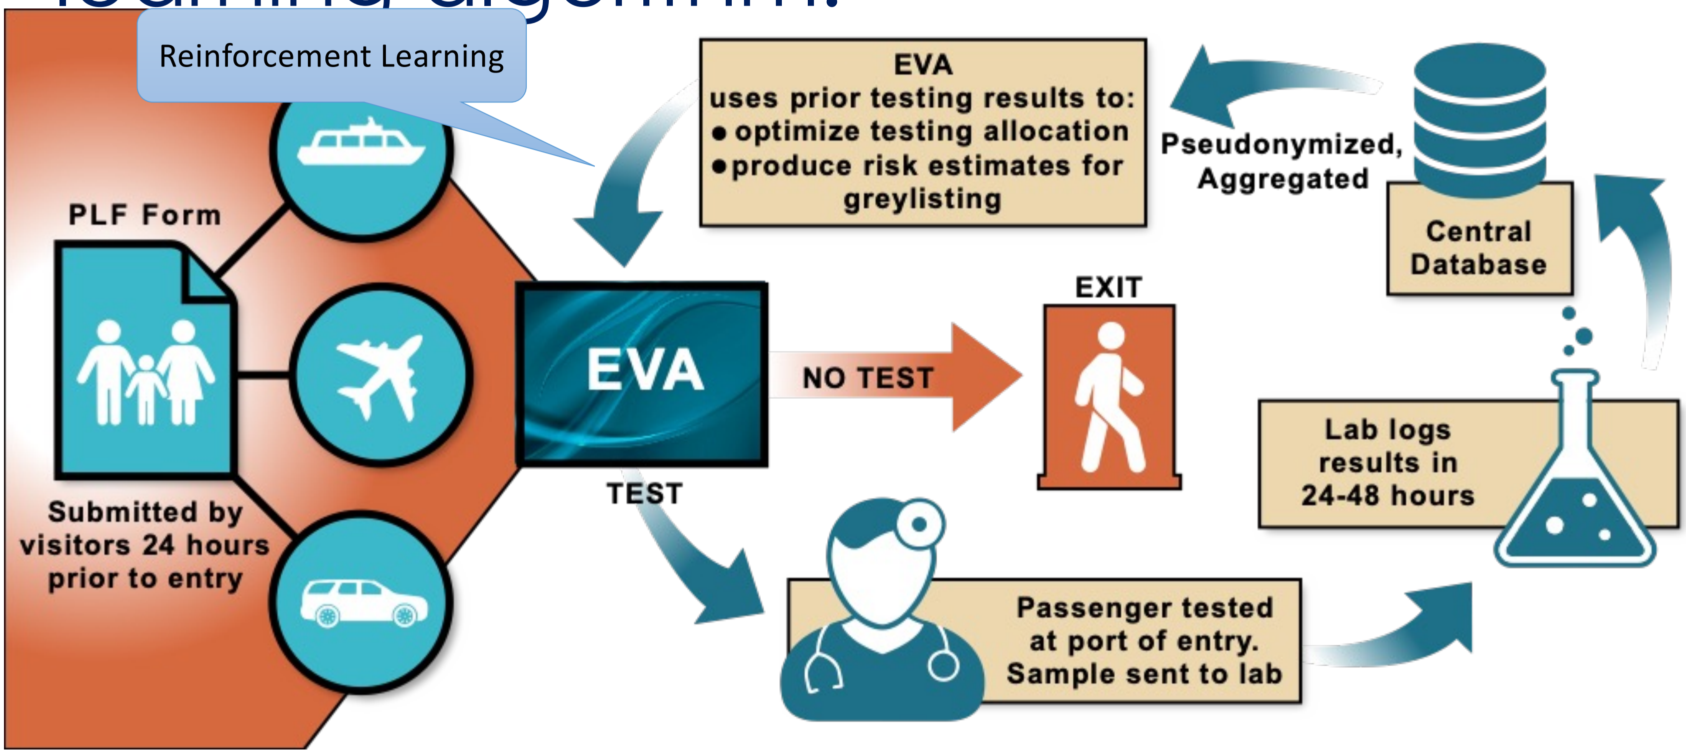

Eva is more than a reinforcement learning algorithm!

GDPR Compliance & EU Support

Digital technologies and data have a valuable role to play in combating the pandemic. Mobile applications could bolster contact-tracing strategies and support public health authorities in monitoring and containing the spread of the virus. Artificial intelligence (AI) and robotics can also help monitoring physical distancing in line with data protection law or facilitating disinfection, especially in places with regular tourism flows. The Commission will deploy through

Reinforcement Learning

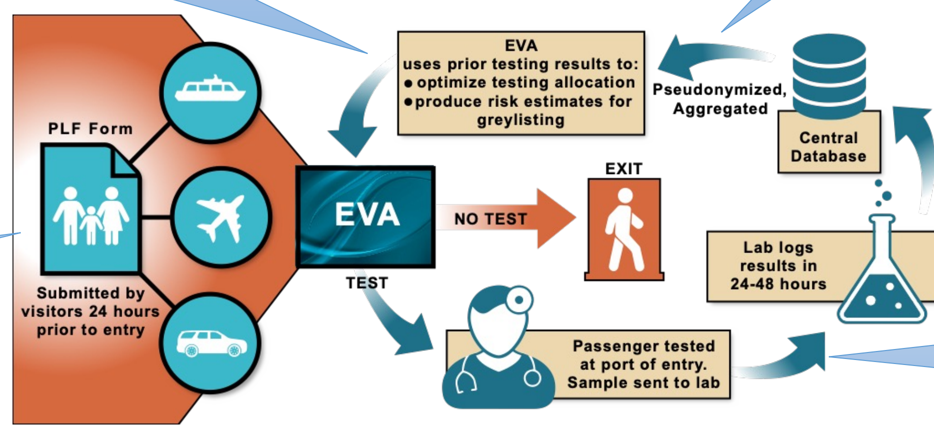

Designing the Passenger Locator Form

- What information is epidemiologically relevant?
- What private information is *too* invasive?

Designing the Testing Supply-Chain

- Which labs should serve which points of entry?
- How transport safely/quickly?

40 Points of Entry (PoE) / 32 Labs

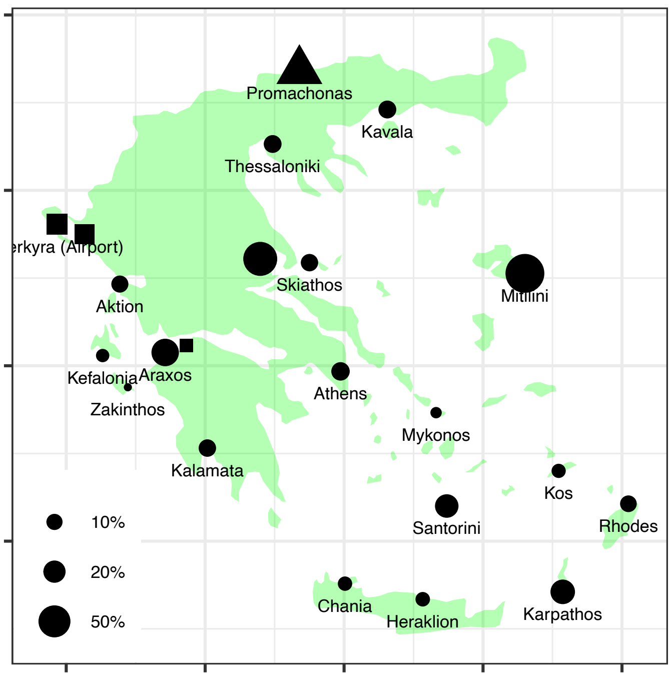

Some PoE and All Labs Omitted for Privacy

- How much testing capacity assign to each PoE?
- Which lab(s) should serve which PoE?

- Single-shot decision / no recourse
  - Model as a mixed-binary optimization

- (Some of the) Constraints
  - Cannot exceed a lab's daily processing capacity
  - Lab/PoE must be close enough for twice-daily trips by logistics teams

Operations

40 points
of entry

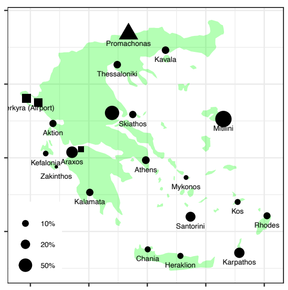

300 medical
professionals
200 border agents

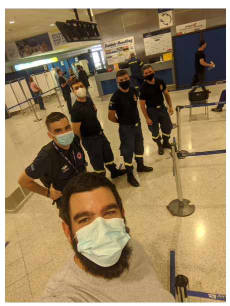

32
laboratories

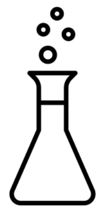

2
teams of
engineers/IT

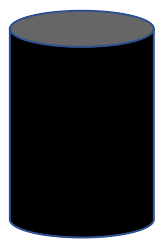

(true prevalence of travelers)

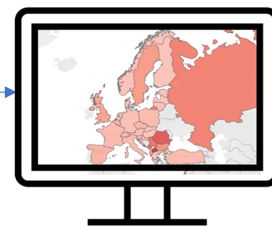

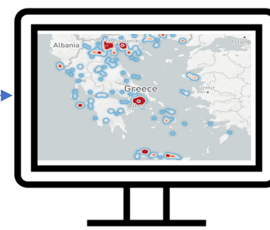

Policy
recommendations

Interplay of Operations Research and Artificial Intelligence

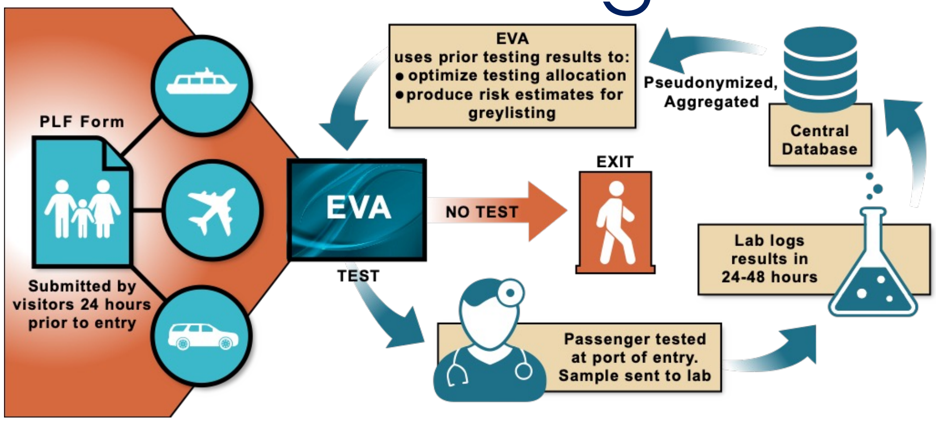

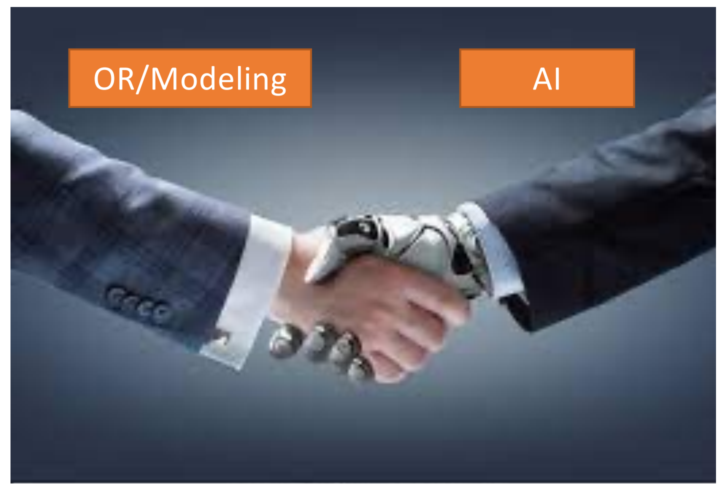

Success of one piece depends on the design of the other…

Expectations and Goals

Real-time data-driven solution to allocate **limited** testing resources

- Recommend **who to test at the border**.
- Identify hot-spots with high confidence to **greylist**

In an ideal world

- "Classic" multiarm bandit problem
  - Estimation: Given current information → $\hat{r}_k(t)$.
  - Allocation:
    - Exploitation: chose arm with largest $\hat{r}_i(t)$.
    - Exploration: Maybe exists j, $r_i(t)$..> $r_j(t)$..

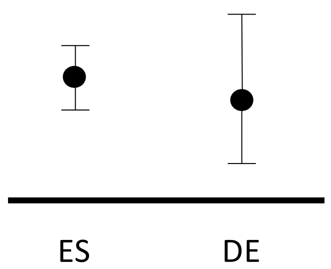

Popular Solutions:

- Upper Confidence Bounds (regret-order-optimal)
- Thomson Sampling (regret-order-optimal)
- Gittins Index (optimal but hard to calculate)

Our setting

- **Types:** countries **(for now)** $k \in \{1, .., K\}$
- Each passenger if tested $X_k \sim B\big(R_k(t)\big)$
- **True positivity** of k-type, $R_k(t)$ **unknown (very) non-stationary** process
- Decision: $N_k(t)$ number of arms to "pull" from each arm
- **Goal:** max $\sum_0^T \sum_{k=0}^K E[T_k(t)],\ T_k(t) \sim Bin(R_k(t), N_k(t))$

Fancy expression for
"more tests to high-risk types"

Combinatorial Constraints

- **I lied** $N_k(t) \rightarrow N_{ke}(t)$
- $N_k(t) = \sum_e N_{ke}(t)$
- Unfortunately:
  - $\sum_k N_{ke}(t) \leq B_e$ (budget)
  - $N_{ke}(t) \leq D_{ke}(\mathrm{t})$ (arrivals)

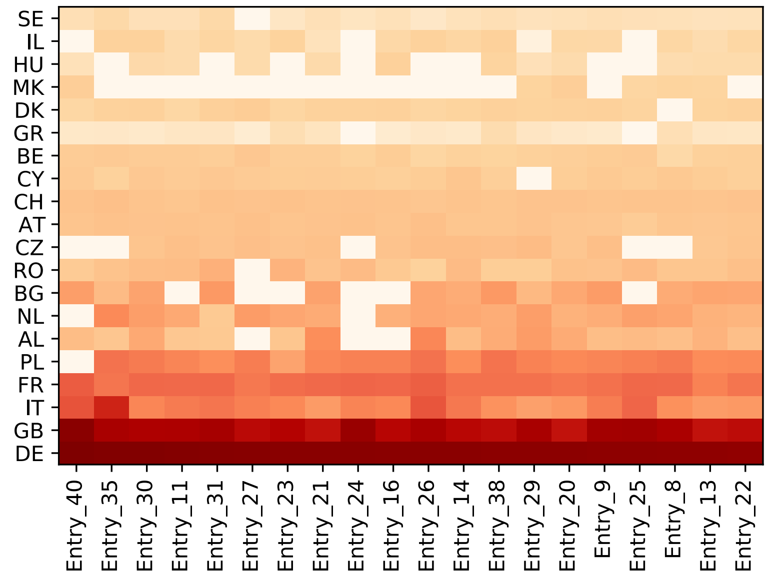

Estimation Challenges

I. Imbalanced data: ~1/1000 test positive
   - How to distinguish 0.7% (risky) vs. 0.1% (safe) arms?
   - Empirical Bayes approach -> fit common prior -> posterior update

II. High-dimensional features: origin city/country/gender etc..
   - Lasso Regression to identify risky subtypes

No technical details

Empirical Bayes approach

- Natural split in groups:
  - **White-listed**
  - **Grey-listed**
  - **Black-listed***
- For each group common prior → $B(\alpha, \beta)$
- Naïve estimators: $P_k/(P_k + N_k)$
- Fit (α, β) → moment matching
- $\boldsymbol{\hat{p}_k(t)}$: posterior based on P_k, N_k

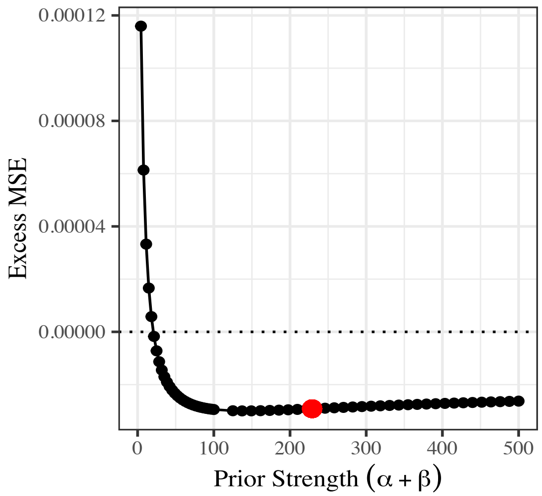

*Defined at the European Union level

Arm Granularity

- **1000s of regions:** cannot allocate 100s of tests to *each* region
  - high variance
- **~100 countries:** miss out on city-specific risk
  - high bias
- Most locations at country-level but isolate very risky cities
- Use LASSO regression to identify a sparse subset of cities that are particularly risky beyond conditioning on country (Bastani & Bayati '15)

Feature Selection

- Segment many cities from a country to a small # of arms

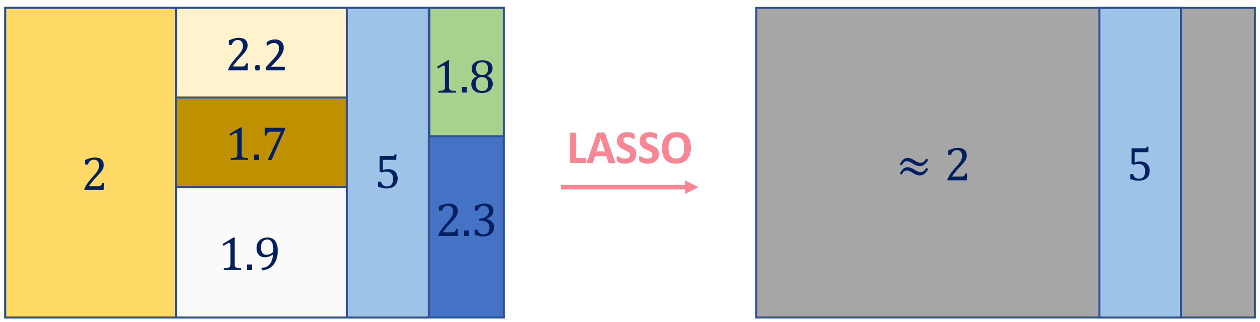

$C_1, C_2, \dots, C_7$ $\quad$ G_1, G_2

Formally: $y_i = \sum_{c=1}^{C} \delta_c \mathbf{1}(c = c_i)\hat{r}_c^{EB} + \sum_f \delta_f \mathbf{1}(f = f_i) + \epsilon_i.$

Back to estimation

Estimation Step

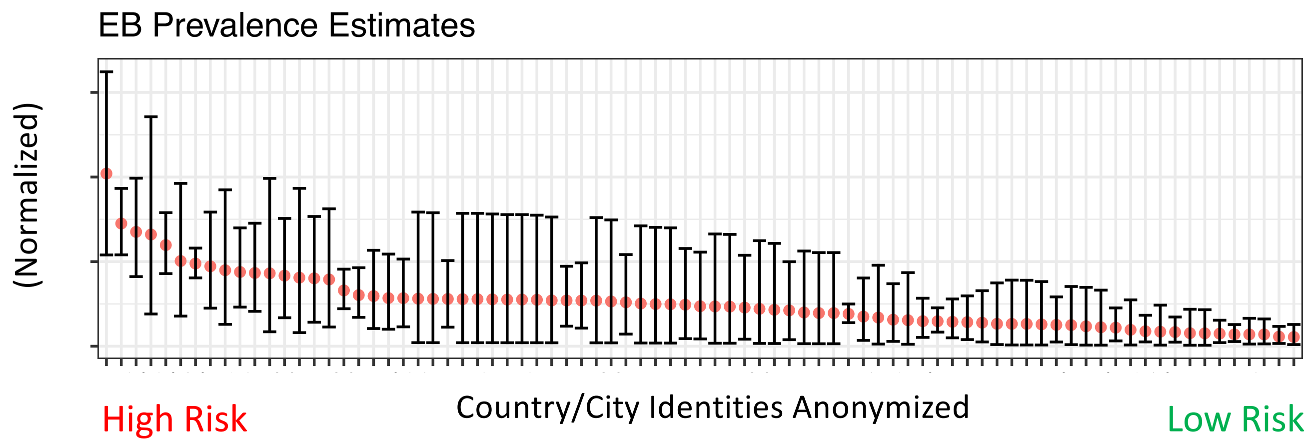

Back to estimation

Allocation Challenges

I. Nonstationary dynamics:
- Discard old data (Luo et al. 2018, Zhao et al. 2020)
- Exponential smoothing (Besbes et al. 2014)

II. Batched decision-making + Delayed feedback:
- decide who to test at the start of each day. 7500 tests at once, 1-3 days to get result.
- Optimistic Gittins Index
- Certainty-Equivalent Pseudo Updates

III. "Combinatorial" Constraints:
- Dynamic matching of indexes to ports.

No technical details

Too forward-looking?

- NO
  - Change ~ 1-2 weeks (Non stationarity)
  - Results delayed ~ 1-3 days
  - "One Step Ahead"

- The optimistic Gittins index (Gutin & Farias 2016) with 1-step lookahead is

$$\lambda = \mathbb{E}[R(y)] + \gamma \cdot \mathbb{E}\left[\left(\lambda - R(y)\right)^{+}\right]$$

- **Intuition:** Bayesian value if we play this arm for one step and then play optimally thereafter
- **Bonus**: simple fixed point equation for Beta distributions

Optimistic Gittins Index

- **Step 1:** Compute λ_i for each arm i
- **Step 2:** Pull arm

$$i^* = \operatorname*{argmax}_i \lambda_i$$

- **Step 3:** Update prior y_i based on observed reward R_t

Within a Batch

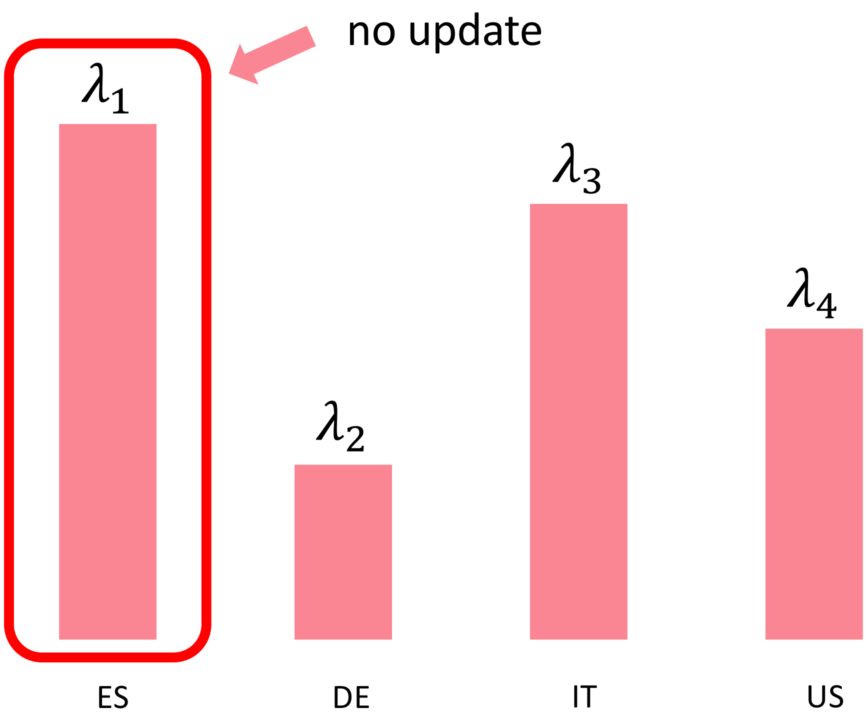

- No new information between pulls
- Over-explore one arm while we are waiting for testing results (2 days)

Certainty Equivalent Update

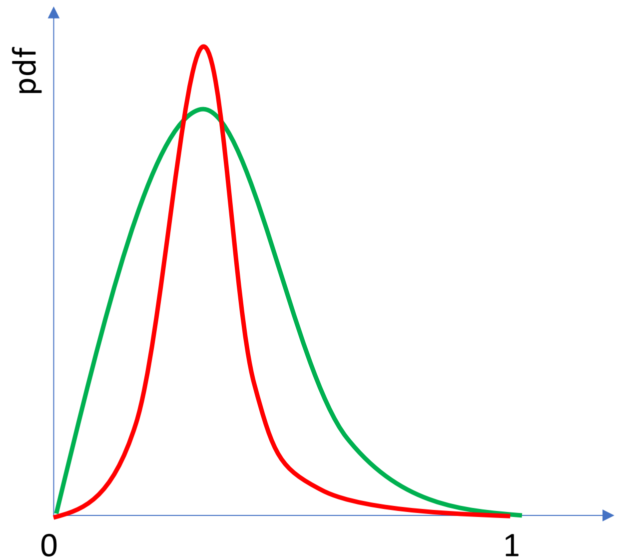

- No new information between pulls
  - but can “simulate” evolution

- Over-explore one arm while we are waiting for testing results (2 days)

Within a Batch

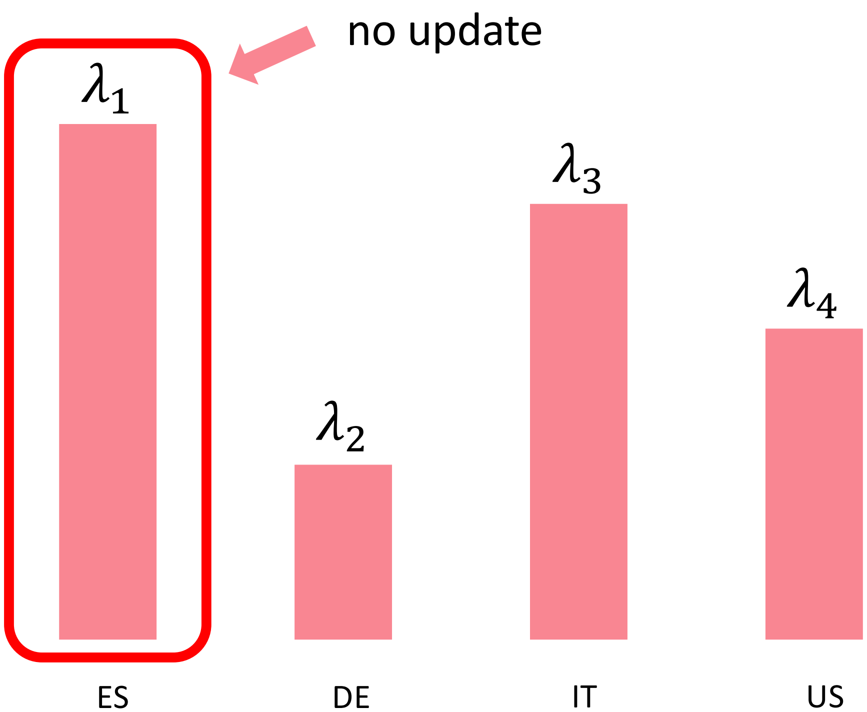

- No new information between pulls
- Over-explore one arm while we are waiting for testing results (2 days)

Within a Batch

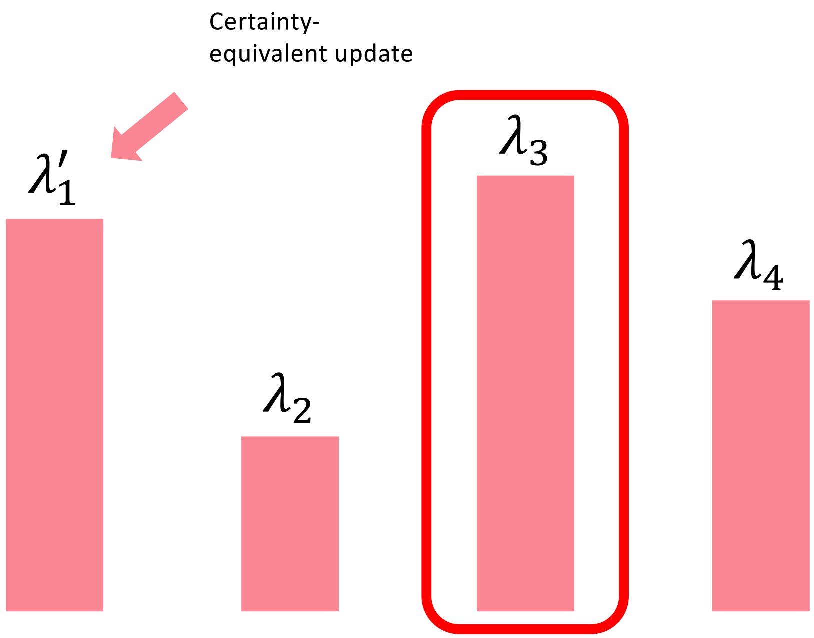

- Allocates estimated # of tests required to resolve uncertainty for arms with high variance
- Remaining tests allocated to arms with high mean reward
- Can simulate tests in "pipeline"

Back to allocation

Combinatorial Constraints

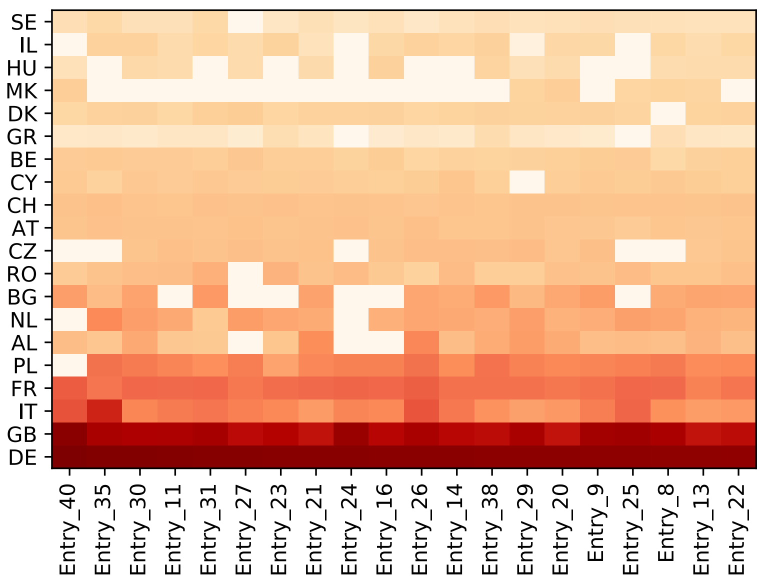

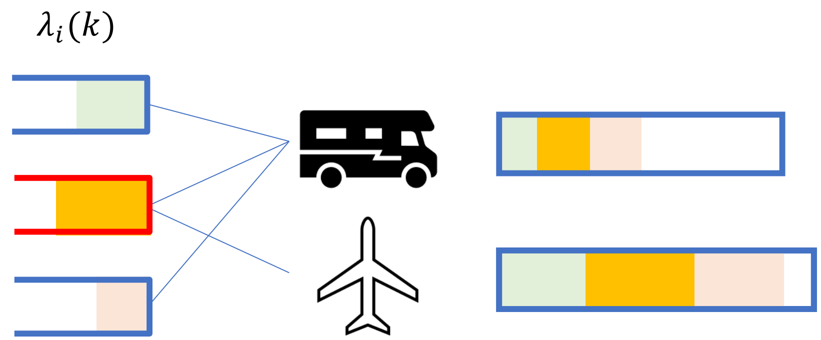

Algorithm

- Inputs:
  - passenger manifest at each port,
  - port-specific testing constraints,
  - historical testing results for each arm
  - certainty equivalent updates for any pending tests

Arm	Pseudo Gittins
1	0.06
2	0.07
3	0.01
⋮	⋮

Arm	Passenger ID	Port + Rem Tests
2	5319	1 (500)
2	2170	3 (50)
1	8562	3 (50)
⋮	⋮	

Algorithm

- Inputs:
  - passenger manifest at each port,
  - port-specific testing constraints,
  - historical testing results for each arm
  - certainty equivalent updates for any pending tests

Arm	Pseudo Gittins
1	0.06
2	0.07
3	0.01
⋮	⋮

Arm	Passenger ID	Port + Rem Tests
2	5319	1 (500)
2	2170	3 (50)
1	8562	3 (50)
⋮	⋮	

Algorithm

- Inputs:
  - passenger manifest at each port,
  - port-specific testing constraints,
  - historical testing results for each arm
  - certainty equivalent updates for any pending tests

Arm	Pseudo Gittins
1	0.06
2	~~0.07~~ 0.05
3	0.01
⋮	⋮

Arm	Passenger ID	Port + Rem Tests
~~2~~	~~5319~~	~~1 (500)~~
2	2170	3 (50)
1	8562	3 (50)
⋮	⋮	

Algorithm

- Inputs:
  - passenger manifest at each port,
  - port-specific testing constraints,
  - historical testing results for each arm
  - certainty equivalent updates for any pending tests

Arm	Pseudo Gittins
1	0.06
2	~~0.07~~ 0.05
3	0.01
⋮	⋮

Arm	Passenger ID	Port + Rem Tests
~~2~~	~~5319~~	~~1 (500)~~
2	2170	3 (50)
1	8562	3 (50)
⋮	⋮	

Back to allocation

Allocations in a Batch

- Exploration & exploitation within a batch

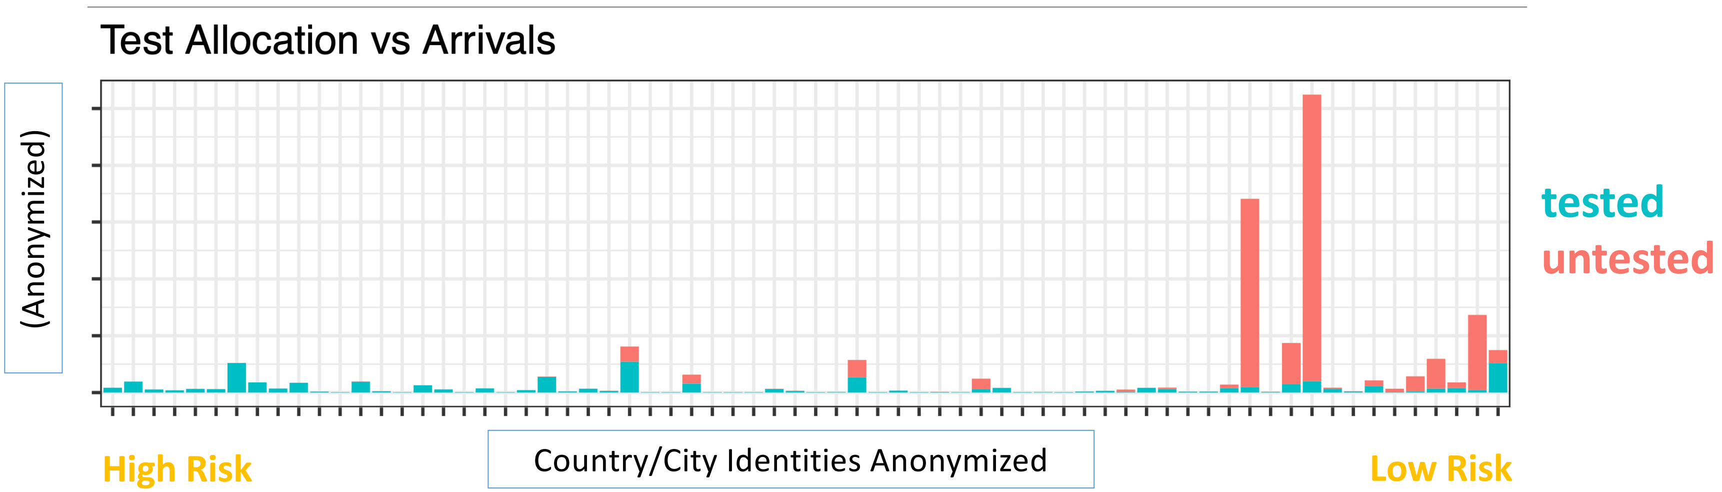

Part II: Evaluation

Benefits: Infections Caught vs. Random Surveillance

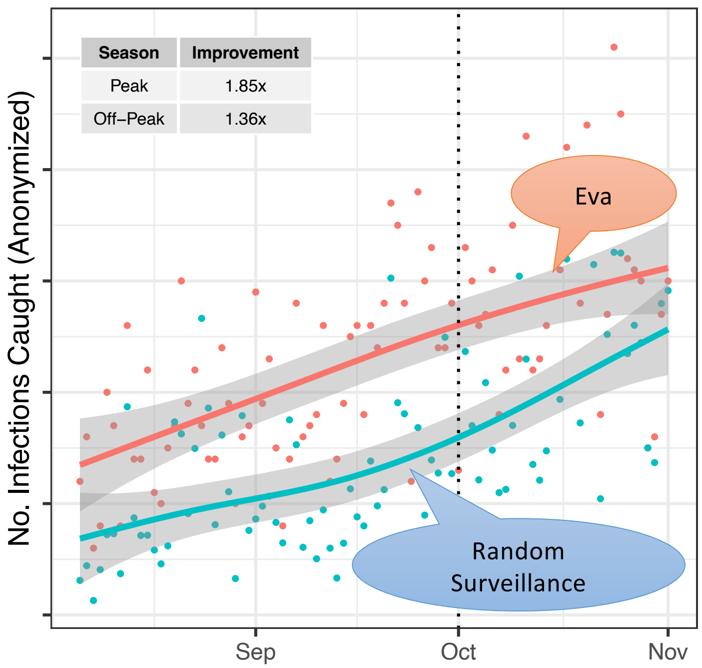

- **Model independent** counterfactual analysis via off-policy learning

- Peak Season (Aug 6 – Sept 30)
  - Random Surveillance identifies 54.1% (±8.7%) of the infections Eva identifies
  - Performance improves as testing is more scarce

Random Surveillance needs 1.85X the testing capacity to achieve same performance!

Improvement vs. Testing Scarcity

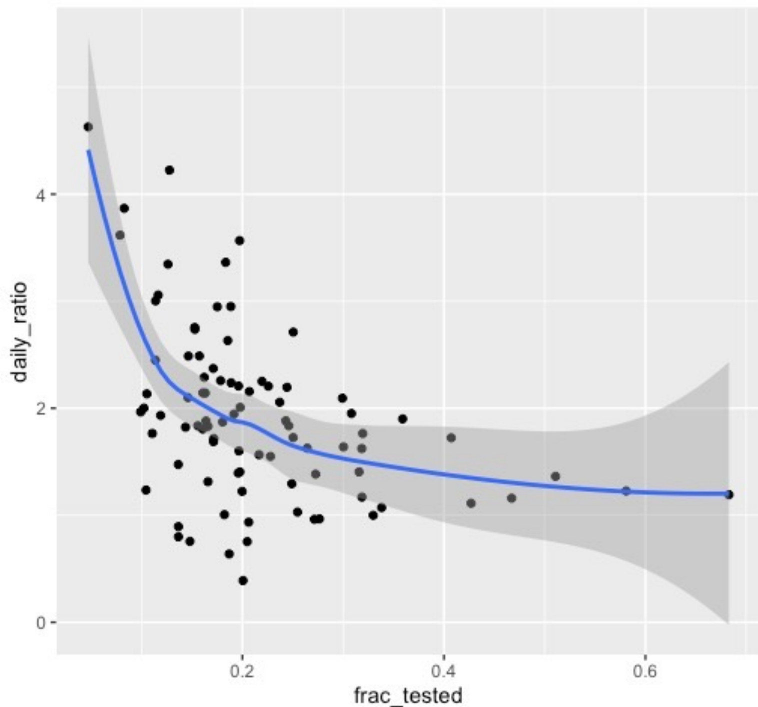

2-3X

when arrivals >> tests

With **7500** at the border

15,000-22,500 tests

More than capacity of the country.

Benefits: Infections Caught vs. “Smart” Surveillance

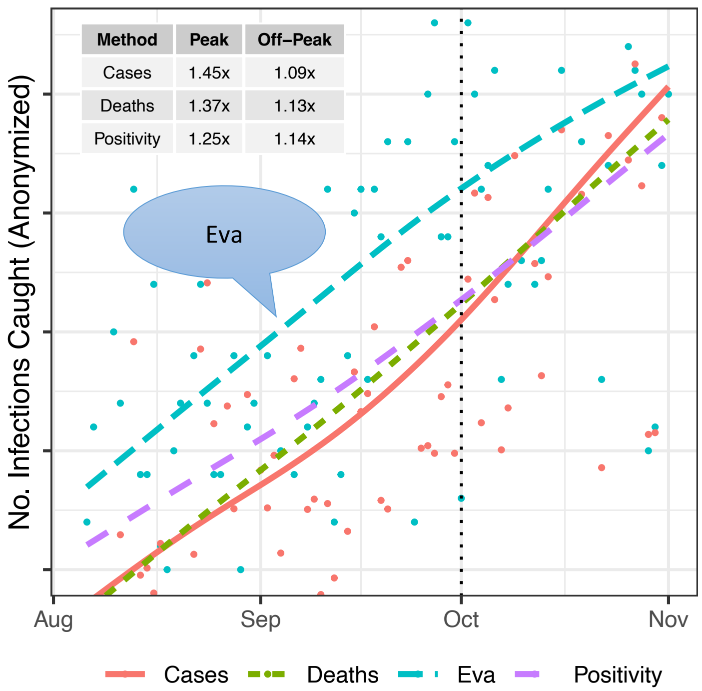

Policies based on Common Epidemiological Data

- A whole HOST of data reliability issues
- Performance ranges from: 69.0% (±9.4%) to 79.7% (±9.3%) of Eva’s Performance

“Smart” Surveillance needs 1.25-1.45X the testing capacity to achieve same performance!

Why? (See paper for details)

- Systematic differences between general population/ asymptomatic traveler population
- Country-specific idiosyncrasies in testing protocols
- Reporting Delays

Ineffectiveness of public data

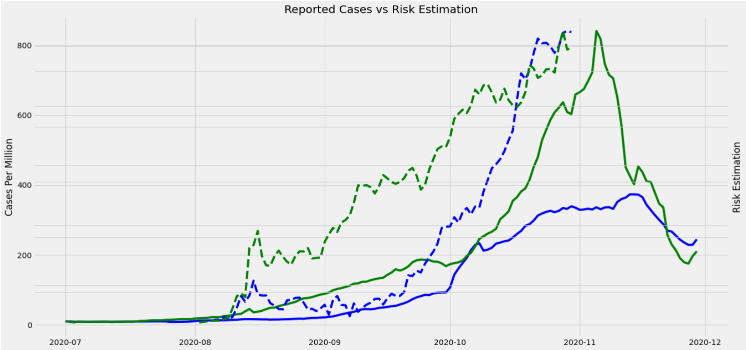

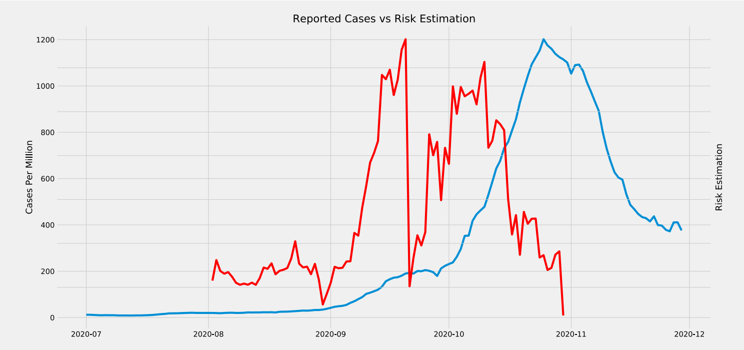

Ineffectiveness of public data

Hypothesis testing

> 0.5% ?

Ineffectiveness of public data

Model	Data Used in Training/Testing
1	14-day average of cases per million, and deaths per million

Hypothesis testing

> 0.5% ?

Ineffectiveness of public data

Model	Data Used in Training/Testing
1	14-day average of cases per million, and deaths per million
2	14-day average of cases per million, deaths per million, tests per thousand and reported positivity rate

Hypothesis testing

> 0.5% ?

Ineffectiveness of public data

Model	Data Used in Training/Testing
1	14-day average of cases per million, and deaths per million
2	14-day average of cases per million, deaths per million, tests per thousand and reported positivity rate
3	14-day time-series of cases per million and deaths per million

Hypothesis testing

> 0.5% ?

Ineffectiveness of public data

Model	Data Used in Training/Testing
1	14-day average of cases per million, and deaths per million
2	14-day average of cases per million, deaths per million, tests per thousand and reported positivity rate
3	14-day time-series of cases per million and deaths per million
4	14-day time-series of cases per million, deaths per million, reported positivity rate, and tests administered per thousand

Hypothesis testing

> 0.5% ?

Ineffectiveness of public data

Model	Data Used in Training/Testing
1	14-day average of cases per million, and deaths per million
2	14-day average of cases per million, deaths per million, tests per thousand and reported positivity rate
3	14-day time-series of cases per million and deaths per million
4	14-day time-series of cases per million, deaths per million, reported positivity rate, and tests administered per thousand
5	14-day time-series of cases per million, deaths per million, positivity rate, tests per thousand, **country fixed effects**

Hypothesis testing

> 0.5% ?

Ineffectiveness of public data

Model	Data Used in Training/Testing
1	14-day average of cases per million, and deaths per million
2	14-day average of cases per million, deaths per million, tests per thousand and reported positivity rate
3	14-day time-series of cases per million and deaths per million
4	14-day time-series of cases per million, deaths per million, reported positivity rate, and tests administered per thousand
5	14-day time-series of cases per million, deaths per million, positivity rate, tests per thousand, **country fixed effects**

Hypothesis testing

> 0.5% ?

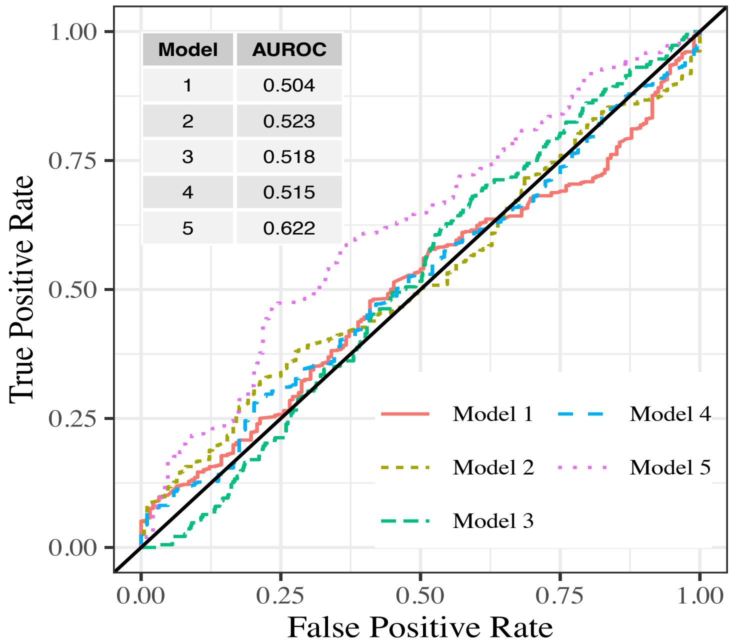

Implications

- To estimate **magnitude** of true prevalence in travelers
  - There is no "one size fits all" rule based on cases, deaths, tests, positivity.
  - Country specific features must be included (e.g. Sun et al. (2020))
  - Model based policies can be problematic (Ahn et al. (2021))

To ensure that the process is manageable and transparent, the proposal focuses on three criteria, namely the 14-day cumulative COVID-19 case notification rate, test positivity rate, and the testing rate. These criteria should then be applied to the different areas, ideally Member States' regions. Only areas with a testing rate of more than 250 COVID-19 tests per 100 000 population should be assessed according to these criteria, to ensure that sufficiently robust data is available.

Model 2: AUC=0.523

Using these criteria, restrictions could be applied, if at all, to regions with a 14-day cumulative COVID-19 case notification of 50 or more and a test positivity rate of 3% or more. Restrictions could be applied to regions where the 14-day cumulative COVID-19 case notification rate is more than 150 per 100 000 population even if the test positivity rate is below 3%. The criteria and thresholds outlined are based on extensive discussions with and data made available by Member States.

Lessons learned the hard way

- <u>"pick your fights":</u> work within practical constraints instead of trying to redesign the organization.
- People **can** be a different kind of smart.
- Don't be arrogant about tasks. Sometime you will have to work on the **menial** tasks.

- Most "real" projects can be politicized. **Avoid** personal **publicity** (at least until the paper is published).
- Don't work on tasks that are not related to your personal goals. It is ok to say **NO.**
- High impact → High stress don't ignore your **emotional health**.

Other Benefits of Eva

Sotirios Tsiodras, MD

- Chief Scientific Advisor, Greek COVID-19 Response
- Member Scientific Advisory Forum European CDC
- President, Greek Infectious Disease Society

Eva used to efficiently allocate scarce testing resources.

Prevalence estimates guided strategy **beyond** targeting

- Reposition mobile testing units within country

Eva's estimates shared centrally with European Union to shape travel policies across continent.

Thank you!

Backup Slides

Simple Network

Let’s think simple:

- 50% of agents have a low degree
- 50% of agents have a high degree

$$\mu_L(t+1) = \mu_L(T) + (1 - \mu_L(t))\Theta(\lambda, t)\ \lambda\ d_L$$

$$\mu_H(t+1) = \mu_H(t) + (1 - \mu_H(t))\Theta(\lambda, t)\ \lambda\ d_H$$

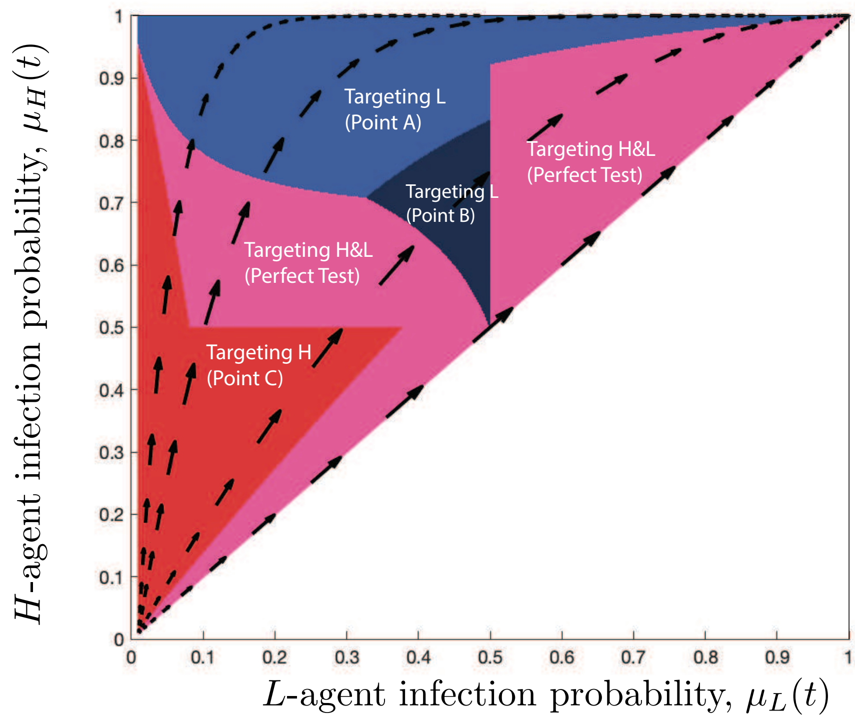

Insights and Results: Lags

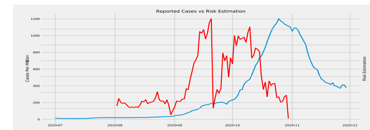

For each country classify whether

$$R_c(t) > median_{summer}(R_c(t'))$$

using cases in the period $[t + \ell - 14, t + \ell]$.

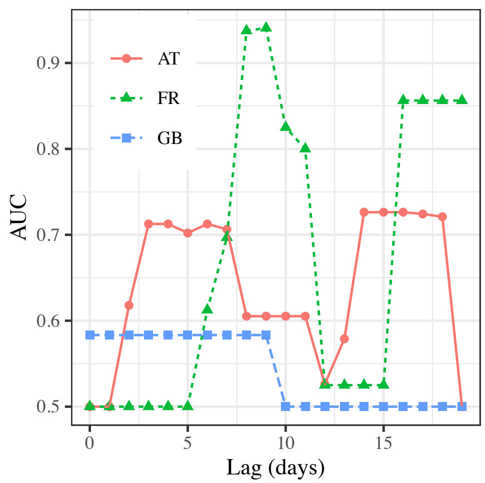

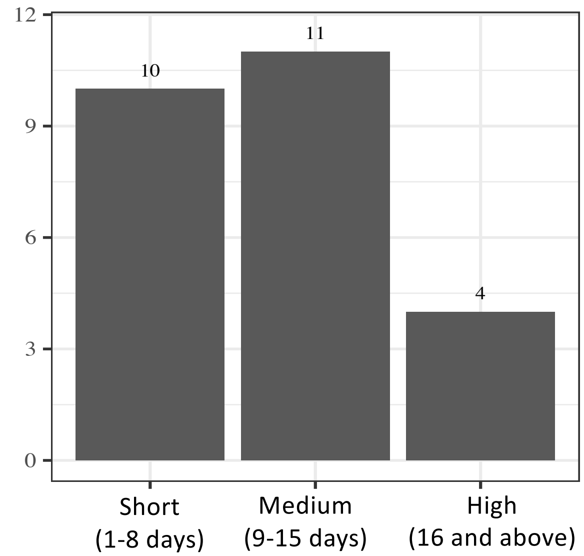

So what?

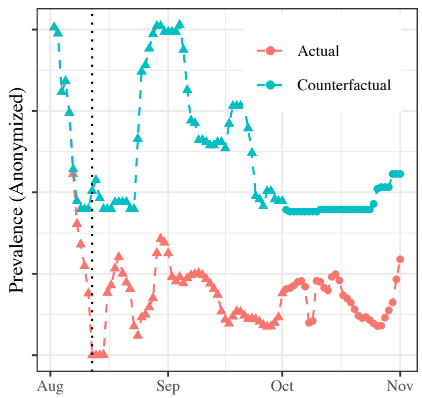

+

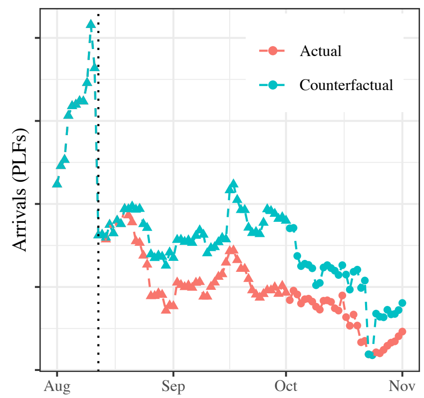

=

1.85x → 2.01x

in peak season

Healthier population

Less arrivals