

---

**Accelerated Article Preview**

---

# Efficient and targeted COVID-19 border testing via reinforcement learning

---

Received: 16 March 2021

---

Accepted: 10 September 2021

---

Accelerated Article Preview Published  
online 22 September 2021

---

Cite this article as: Bastani, H. et al. Efficient and targeted COVID-19 border testing via reinforcement learning. *Nature* <https://doi.org/10.1038/s41586-021-04014-z> (2021).

---

Hamsa Bastani, Kimon Drakopoulos, Vishal Gupta, Jon Vlachogiannis, Christos Hadjicristodoulou, Pagona Lagiou, Gkikas Magiorkinis, Dimitrios Paraskevis & Sotirios Tsiodras

---

This is a PDF file of a peer-reviewed paper that has been accepted for publication. Although unedited, the content has been subjected to preliminary formatting. Nature is providing this early version of the typeset paper as a service to our authors and readers. The text and figures will undergo copyediting and a proof review before the paper is published in its final form. Please note that during the production process errors may be discovered which could affect the content, and all legal disclaimers apply.

# Efficient and targeted COVID-19 border testing via reinforcement learning

<https://doi.org/10.1038/s41586-021-04014-z>

Received: 16 March 2021

Accepted: 10 September 2021

Published online: 22 September 2021

Hamsa Bastani<sup>1,7</sup>, Kimon Drakopoulos<sup>2,7✉</sup>, Vishal Gupta<sup>2,7</sup>, Jon Vlachogiannis<sup>3</sup>, Christos Hadjicristodoulou<sup>4</sup>, Pagona Lagiou<sup>5</sup>, Gkikas Magiorkinis<sup>5</sup>, Dimitrios Paraskevis<sup>5</sup> & Sotirios Tsiodras<sup>6</sup>

Throughout the COVID-19 pandemic, countries relied on a variety of ad-hoc border control protocols to allow for non-essential travel while safeguarding public health: from quarantining all travellers to restricting entry from select nations based on population-level epidemiological metrics such as cases, deaths or testing positivity rates<sup>1,2</sup>. Here we report the design and performance of a reinforcement learning system, nicknamed ‘Eva’. In the summer of 2020, Eva was deployed across all Greek borders to limit the influx of asymptomatic travellers infected with SARS-CoV-2, and to inform border policies through real-time estimates of COVID-19 prevalence. In contrast to country-wide protocols, Eva allocated Greece’s limited testing resources based upon incoming travellers’ demographic information and testing results from previous travellers. By comparing Eva’s performance against modelled counterfactual scenarios, we show that Eva identified 1.85 times as many asymptomatic, infected travellers as random surveillance testing, with up to 2–4 times as many during peak travel, and 1.25–1.45 times as many asymptomatic, infected travellers as testing policies that only utilize epidemiological metrics. We demonstrate that this latter benefit arises, at least partially, because population-level epidemiological metrics had limited predictive value for the actual prevalence of SARS-CoV-2 among asymptomatic travellers and exhibited strong country-specific idiosyncrasies in the summer of 2020. Our results raise serious concerns on the effectiveness of country-agnostic internationally proposed border control policies<sup>3</sup> that are based on population-level epidemiological metrics. Instead, our work represents a successful example of the potential of reinforcement learning and real-time data for safeguarding public health.

In the first wave of the pandemic, many countries restricted non-essential travel to mitigate the spread of SARS-CoV-2. The restrictions crippled most tourist economies, with estimated losses of 1 trillion USD among European countries and 19 million jobs<sup>3</sup>. As conditions improved from April to July, countries sought to partially lift these restrictions, not only for tourists, but also for the flow of goods and labor.

Different countries adopted different border screening protocols, typically based upon the origin country of the traveler. Despite their variety, we group the protocols used in early summer 2020 into 4 broad types:

- Allowing unrestricted travel from designated “white-list” countries.
- Requiring travelers from designated “grey-listed” countries to provide proof of a negative RT-PCR test before arrival.
- Requiring all travelers from designated “red-listed” countries to quarantine upon arrival.

- Forbidding any non-essential travel from designated “black-listed” countries.

Most nations employed a combination of all four strategies. However, the choice of which “color” to assign to a country differed across nations. For example, as of July 1<sup>st</sup>, 2020, Spain designated the countries specified in<sup>1</sup> as white-listed, Croatia designated these countries as grey-listed or red-listed.

To the best of our knowledge, in all European nations except Greece, the above “color designations” were entirely based on population-level epidemiological metrics (e.g., see<sup>1,2</sup>) such as cases per capita, deaths per capita, and/or positivity rates that were available in the public domain<sup>4–6</sup>. (An exception is the UK, which engaged in small-scale testing at select airports that *may* have informed their policies.) However such metrics are imperfect due to underreporting<sup>7</sup>, symptomatic population biases<sup>8–10</sup> and reporting delays.

<sup>1</sup>Department of Operations, Information and Decisions, Wharton School, University of Pennsylvania, Philadelphia, PA, USA. <sup>2</sup>Department of Data Sciences and Operations, Marshall School of Business, University of Southern California, Los Angeles, CA, USA. <sup>3</sup>AgentRisk, Los Angeles, CA, USA. <sup>4</sup>Department of Hygiene and Epidemiology, University of Thessaly, Thessaly, Greece.

<sup>5</sup>Department of Hygiene, Epidemiology and Medical Statistics, School of Medicine, National and Kapodistrian University of Athens, Athens, Greece. <sup>6</sup>Department of Internal Medicine, Attikon University Hospital, Medical School, National and Kapodistrian University of Athens, Athens, Greece. <sup>7</sup>These authors contributed equally: Hamsa Bastani, Kimon Drakopoulos, Vishal Gupta.

✉e-mail: drakopou@marshall.usc.edu

These drawbacks motivated our design and nation-wide deployment of Eva: the first fully algorithmic, real-time, reinforcement learning system for targeted COVID-19 screening with the dual goals of identifying asymptomatic, infected travelers and providing real-time information to policymakers for downstream decisions.

## The Eva System: Overview

Eva as presented here was deployed across all 40 points of entry to Greece, including airports, land crossings, and seaports from August 6<sup>th</sup> to November 1<sup>st</sup>. Fig. 1 schematically illustrates its operation; Fig. 7 in Methods provides a more detailed schematic of Eva's architecture and data flow.

### 1. Passenger Locator Form (PLF)

All travelers must complete a PLF (one per household) at least 24 hours prior to arrival, containing (among other data) information on their origin country, demographics, point and date of entry<sup>11</sup>, describes the exact fields and how these sensitive data were handled securely.

### 2. Estimating Prevalence among Traveler Types

We estimate traveler-specific COVID-19 prevalence using recent testing results from previous travelers through Eva. Prevalence estimation entailed two steps: First, we leverage LASSO regression from high-dimensional statistics<sup>12</sup> to adaptively extract a minimal set of discrete, interpretable traveler *types* based on their demographic features (country, region, age, gender); these types are updated on a weekly basis using recent testing results. Second, we use an empirical Bayes method to estimate each type's prevalence daily. Empirical Bayes has previously been used in the epidemiological literature to estimate prevalence across many populations<sup>13,14</sup>. In our setting, COVID-19 prevalence is generally low (e.g., ~2 in 1000), and arrival rates differ substantively across countries. Combined, these features cause our testing data to be both imbalanced (few positive cases among those tested) and sparse (few arrivals from certain countries). Our empirical Bayes method seamlessly handles both challenges. Estimation details are provided in Section 2.2 of Methods.

### 3. Allocating Scarce Tests

Leveraging these prevalence estimates, Eva targets a subset of travelers for (group) PCR testing upon arrival based solely on their type, but no other personal information. The Greek National COVID-19 Committee of Experts approved group (Dorfman) testing<sup>15</sup> in groups of 5 but eschewed larger groups and rapid testing due to concerns over testing accuracy.

Eva's targeting must respect various port-level budget and resource constraints that reflect Greece's testing supply chain, which included 400 health workers staffing 40 points of entry, 32 laboratories across the country, and delivery logistics for biological samples. These constraints were (exogenously) defined and adjusted throughout the summer by the General Secretariat of Public Health.

The *testing allocation* decision is entirely algorithmic and balances two objectives: First, given current information, Eva seeks to maximize the number of infected asymptomatic travelers identified (*exploitation*). Second, Eva strategically allocates some tests to traveler types for which it does not currently have precise estimates in order to better learn their prevalence (*exploration*). *This is a crucial feedback step*. Today's allocations will determine the available data in Step 2 above when determining future prevalence estimates. Hence, if Eva simply (greedily) sought to allocate tests to types that currently had high prevalence, then, in a few days, it would not have any recent testing data about many other types that had moderate prevalence. Since COVID-19 prevalence can spike quickly and unexpectedly, this would leave a "blind spot" for the algorithm and pose a serious public health risk. Such allocation problems can be formulated as multi-armed

bandits<sup>16-19</sup> – which are widely studied within the reinforcement learning literature – and have been used in numerous applications such as mobile health<sup>20</sup>, clinical trial design<sup>21</sup>, online advertising<sup>22</sup>, and recommender systems<sup>23</sup>.

Our application is a nonstationary<sup>24,25</sup>, contextual<sup>26</sup>, batched bandit problem with delayed feedback<sup>27,28</sup> and constraints<sup>29</sup>. Although these features have been studied in isolation, their combination and practical implementation poses unique challenges. One such challenge is accounting for information from "pipeline" tests (allocated tests whose results have not yet been returned); we introduce a novel algorithmic technique of *certainty-equivalent* updates to model information we *expect* to receive from these tests, allowing us to effectively balance exploration and exploitation in nonstationary, batched settings. To improve interpretability, we build on the optimistic gittins index for multi-armed bandits;<sup>30</sup> each type is associated with a deterministic index that represents its current "risk score", incorporating both its estimated prevalence and uncertainty. Algorithm details are provided in Section 2.3 of Methods.

### 4. Grey-Listing Recommendations

Eva's prevalence estimates are also used to recommend particularly risky countries to be grey-listed, in conjunction with the Greek COVID-19 taskforce and the Presidency of the Government. Grey-listing a country entails a tradeoff: requiring a PCR test reduces the prevalence among incoming travelers, however, it also reduces non-essential travel *significantly* (approximately 39%, c.f. Sec. 3.2 of Methods), because of the relative difficulty/expense in obtaining PCR tests in summer 2020. Hence, Eva recommends grey-listing a country only when necessary to keep the daily flow of (uncaught) infected travelers at a sufficiently low level to avoid overwhelming contact-tracing teams<sup>31</sup>. Ten countries were grey-listed over the summer of 2020 (see Sec. 4 of Methods).

Unlike testing decisions, our grey-listing decisions were not fully algorithmic, but instead involved human input. Indeed, while in theory, one might determine an "optimal" cut-off for grey-listing to balance infected arrivals and reduced travel, in practice it is difficult to elicit such preferences from decision-makers directly. Rather, they preferred to retain some flexibility in grey-listing to consider other factors in their decisions.

### 5. Closing the Loop

Results from tests performed in Step 3 are logged within 24-48 hours, and then used to update the prevalence estimates in Step 2.

To give a sense of scale, during peak season (August and September), Eva processed 41,830 ( $\pm 12,784$ ) PLFs each day, and 16.7% ( $\pm 4.8\%$ ) of arriving households were tested each day.

## Results

### Value of Targeted Testing: Cases Identified

We first present the number of asymptomatic, infected travelers caught by Eva relative to random surveillance testing, i.e., where every arrival at a port of entry is equally likely to be tested. Random surveillance testing was Greece's initial proposal and is very common, partly because it requires no information infrastructure to implement. However, we find that such an approach comes at a significant cost to performance (and therefore public health).

We perform counterfactual analysis using inverse propensity weighting (IPW)<sup>32,33</sup>, which provides a *model-agnostic, unbiased* estimate of the performance of random testing.

During the peak tourist season, we estimate that random surveillance testing would have identified 54.1% ( $\pm 8.7\%$ ) of the infected travelers that Eva identified. (For anonymity, averages and standard deviations are scaled by a (fixed) constant, which we have taken without loss of generality to be the actual number of infections identified by Eva in the same period for ease of comparison.)

In other words, to achieve the same effectiveness as Eva, random testing would have required 85% *more* tests at each point of entry, a substantive supply chain investment. In October, when arrival rates dropped, the relative performance of surveillance testing improved to 73.4% ( $\pm 11.0\%$ ) (see Fig. 2). This difference is largely explained by the changing relative scarcity of testing resources (see Fig. 3). As arrivals dropped, the fraction of arrivals tested increased, thereby reducing the value of targeted testing. Said differently, Eva's targeting is most effective when tests are scarce. In the extreme case of testing 100% of arrivals, targeted testing offers no value since both random and targeted testing policies test everyone. See Sec. 3.1 of Methods for details.

### Value of Reinforcement Learning: Cases Identified

We now compare to policies that require similar infrastructure as Eva, namely PLF data, but instead target testing based on population-level epidemiological metrics (e.g., as proposed by the EU<sup>2</sup>) rather than reinforcement learning. The financial investments of such approaches are similar to those of Eva, and we show these policies identify fewer cases. (Sec. 3.2.3 of Methods highlights additional drawbacks of these policies, including poor data reliability and a mismatch in prevalence between the general population and asymptomatic traveler population.)

We consider three separate policies that test passengers with probability proportional to either (i) cases per capita, (ii) deaths per capita, or (iii) positivity rates for the passenger's country of origin<sup>4-6</sup>, while respecting port budgets and arrival constraints. We again use IPW to estimate counterfactual performance (see Fig. 4).

During the peak tourist season (August, September), we found that policies based on cases, deaths and positivity rates identified 69.0% ( $\pm 9.4\%$ ), 72.7% ( $\pm 10.6\%$ ), and 79.7% ( $\pm 9.3\%$ ) respectively of the infected travelers that Eva identified per test. In other words, Eva identified 1.25x – 1.45x more infections with the same testing budget and similar PLF infrastructure. In October, when arrival rates dropped, the relative performance of counterfactual policies based on cases, deaths and positivity rates improved to 91.5% ( $\pm 11.7\%$ ), 88.8% ( $\pm 10.5\%$ ) and 87.1% ( $\pm 10.4\%$ ) respectively. Like our results in the previous section, we find that the value of smart targeting is larger when testing resources are scarcer. In fact, Eva's relative improvement over these policies was highest in the second half of the peak season (when infection rates were much higher and testing resources were scarcer). See Sec. 3.2 of Methods for details.

Sec. 4 of Methods discusses possible reasons underlying the poor performance of simple policies based on population-level epidemiological metrics, including reporting delays and systematic differences between the general and asymptomatic traveler populations.

### Poor Predictive Power of Population-Level Epidemiological Metrics

Given the poor performance of simple policies based on population-level epidemiological metrics, a natural question is whether more sophisticated functions of these metrics would perform better. Although it is difficult to eliminate this possibility, we argue this is likely not the case by studying a related question: "To what extent can population-level epidemiological metrics be used to predict COVID-19 prevalence among asymptomatic travelers as measured by Eva?" To the best of our knowledge, this is the first study of this kind. Surprisingly, our findings suggest that widely used epidemiological data are generally *ineffective* in predicting the actual prevalence of COVID-19 among asymptomatic travelers (the group of interest for border control policies).

Specifically, we examine the extent to which these data can be used to classify a country as high-risk (more than 0.5% prevalence) or low-risk (less than 0.5% prevalence); such a classification informs whether a country should be grey- or black-listed. (A cutoff of 0.5% was typical for

initiating grey-listing discussions with the Greek COVID-19 taskforce, but our results are qualitatively similar across a range of cutoffs.) We compute the true label for a country at each point in time based on Eva's (real-time) estimates. We then train several models using a Gradient Boosted Machine (GBM)<sup>34</sup> on different subsets of covariates derived from the 14-day time series of cases per capita, deaths per capita, testing rates per capita, and testing positivity rates. Fig. 5 summarizes their predictive accuracy; we obtained similar results for other state-of-the-art machine learning algorithms.

Note that a random model that uses no data has an AUROC of 0.5. Thus, Models 1-4 offer essentially no predictive value, suggesting that these population-level epidemiological data are not informative of prevalence among asymptomatic travelers.

Model 5, which additionally uses country-level fixed effects, offers some improvement. These fixed effects collectively model country-specific idiosyncrasies representing aspects of their testing strategies, social distancing protocols and other non-pharmaceutical interventions that are *unobserved* in the public, epidemiological data. The improvements of Model 5 suggests that these unobserved drivers are critical to distinguishing high- and low-risk countries.

Overall, this analysis not only raises concerns about travel protocols proposed by the EU<sup>2</sup> based solely upon widely used epidemiological metrics, but also about *any* protocol that treats all countries symmetrically. Indeed, the idiosyncratic effects of Model 5 suggest that the thresholds for deciding whether COVID-19 prevalence in travelers from Country A is spiking may differ significantly from that of Country B. See Section 4.1 for details.

In Section 4.3 of Methods, we also study the information *delay* between a country's publicly reported cases (the most common metric) and prevalence among asymptomatic travelers from that country. We expect a lag because of the time taken for symptoms to manifest, and reporting delays induced by poor infrastructure. We find a modal delay of 9 days.

### Value of Grey-Listing: Cases Prevented

Eva's measurements of COVID-19 prevalence were also used to provide early warnings for high-risk regions, in response to which Greece adjusted travel protocols by grey-listing these nations. We estimate that Eva prevented an additional 6.7% ( $\pm 1.2\%$ ) infected travelers from entering the country through its early grey-listing decisions in the peak season; results in the off-peak season are similar. For privacy, we have expressed the benefit relative to the number of infected travelers identified by Eva. See Sec. 5 of Methods for details.

### Conclusions: Lessons Learned from Deployment and Design Principles

To the best of our knowledge, Eva was the first large-scale data-driven system that was designed and deployed during the COVID-19 crisis. Leading up to and throughout deployment, we met twice a week with the COVID-19 Executive Committee of Greece, an interdisciplinary team of scientists and policymakers. Through those meetings, we gleaned several lessons that shaped Eva's design and contributed to its success.

*Design the algorithm around data minimization.* Data minimization (i.e., requesting the minimum required information for a task), is a fundamental tenet of data privacy and the General Data Protection Regulation (GDPR). We met with lawyers, epidemiologists, and policymakers *before* designing the algorithm to determine what data and granularity may legally and ethically be solicited by the PLF. Data minimization naturally entails a tradeoff between privacy and effectiveness. We limited requests to features thought to be predictive based on best available research at the time (origin, age and gender<sup>35,36</sup>), but omitted potentially informative but invasive features (e.g., occupation). We further designed our empirical Bayes estimation strategy around these data limitations.

**Prioritize interpretability.** For all parties to evaluate and trust the recommendations of a system, the system must provide transparent reasoning. An example from our deployment was the need to communicate the rationale for “exploration” tests, i.e., tests for types with moderate but very uncertain prevalence estimates). Such tests may seem wasteful. Our choice of empirical Bayes allowed us to easily communicate that types with large confidence intervals may have significantly higher risk than their point estimate suggests, and thus require some tests to resolve uncertainty; see, e.g., Figs. 9 and 11 in Methods, which were featured on policymakers’ dashboards.

A second example was our choice to use gittins indices, which provide a simple, deterministic risk metric for each type that incorporates both estimated prevalence and corresponding uncertainty, driving intuitive test allocations. In contrast, using UCB or Thompson Sampling with logistic regression<sup>37,38</sup> would have made it more difficult to visualize uncertainty (a high-dimensional ellipsoid or posterior distribution) and test allocations would depend on this uncertainty through an opaque computation (a high-dimensional projection or stochastic sampling).

This transparency fostered trust across ministries of the Greek Government using our estimates to inform downstream policy making, including targeting contact-tracing teams, staffing of mobile testing units, and adjusting social distancing measures.

**Design for flexibility.** Finally, since these systems require substantial financial and technical investment, they need to be flexible to accommodate unexpected changes. We designed Eva in a modular manner disassociating type extraction, estimation, and test allocation. Consequently, one module can easily be updated without altering the remaining modules. For example, had vaccine distribution begun summer 2020, we could define new types based on passengers’ vaccination status without altering our procedure for prevalence estimates or test allocation. Similarly, if rapid testing were approved, our allocation mechanism could be updated to neglect delayed feedback without affecting other components. This flexibility promotes longevity, since it is easier to get stakeholder buy-in for small adjustments to an existing system than for a substantively new approach.

## Online content

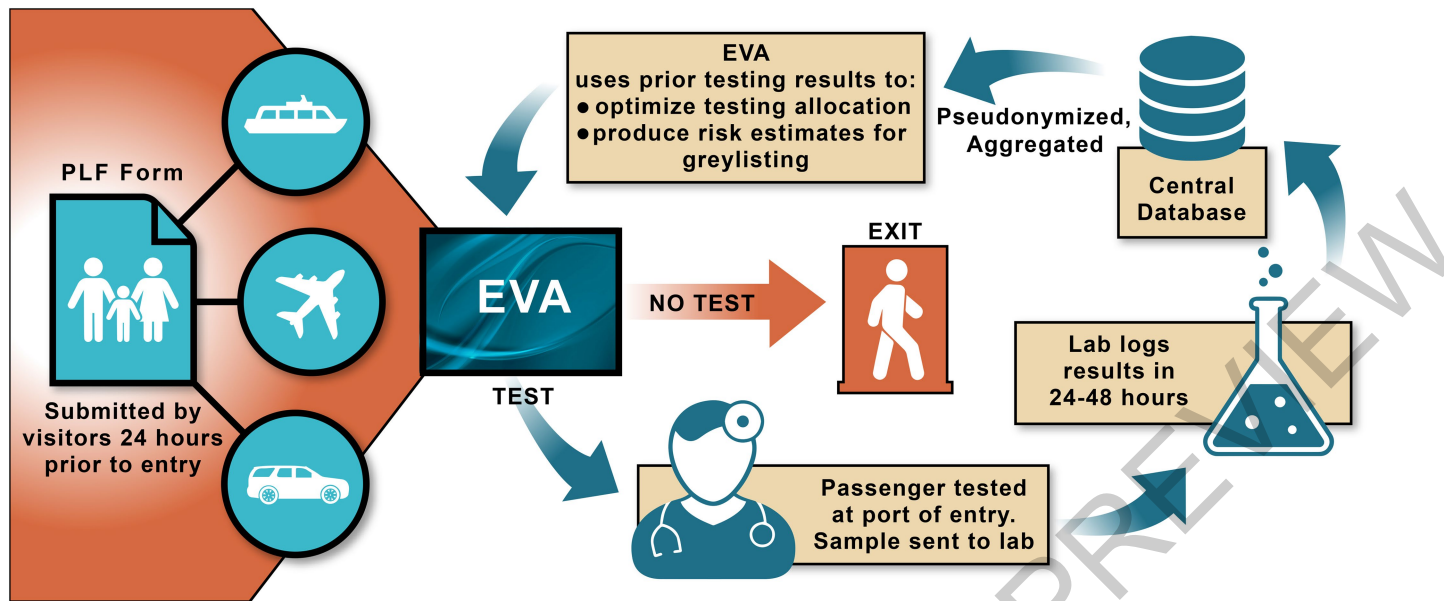
Any methods, additional references, Nature Research reporting summaries, source data, extended data, supplementary information, acknowledgements, peer review information; details of author contributions and competing interests; and statements of data and code availability are available at <https://doi.org/10.1038/s41586-021-04014-z>.

1. Council recommendation on the temporary restriction on non-essential travel into the EU and the possible lifting of such restriction, Brussels, 30 June 2020: [https://www.consilium.europa.eu/media/47592/st\\_9208\\_2020\\_init\\_en.pdf](https://www.consilium.europa.eu/media/47592/st_9208_2020_init_en.pdf).
2. Draft Council Recommendation on a coordinated approach to the restriction of free movement in response to the COVID-19 pandemic, Brussels, 12 October 2020: General Secretariat of the Council.
3. “World Travel and Tourism Council,” November 2020. [Online]. Available: Council <https://wtcc.org/Research/Economic-Impact/Recovery-Scenarios>.
4. J. Hasell, E. Mathieu, D. Beltekian, B. a. G. C. a. O.-O. E. Macdonald, M. Roser and H. Ritchie, “A cross-country database of COVID-19 testing,” *Scientific data*, vol. 7, no. 1, pp. 1-7, 2020.
5. M. Roser, H. Ritchie, E. Ortiz-Ospina and J. Hasell, “Coronavirus Pandemic (COVID-19),” OurWorldInData.org, 2020. [Online]. Available: <https://ourworldindata.org/coronavirus>.
6. E. Dong, H. Du, and L. Gardner, “An interactive web-based dashboard to track COVID-19 in real time,” *The Lancet infectious diseases*, vol. 20, no. 5, pp. 533-534, 2020.
7. S. L. Wu, A. N. Mertens, Y. S. Crider, A. Nguyen, N. N. Pokpongkiat, S. Djajadi, A. Seth, M. S. Hsiang, J. M. J. Colford, A. Reingold, B. F. Arnold, A. Hubbard and J. Benjamin-Chung, “Substantial underestimation of SARS-CoV-2 infection in the United States,” *Nature Communications*, 2020.

8. P. Zhao, L. Zhang, Y. Jiang and Z.-H. Zhou, “A simple approach for non-stationary linear bandits,” *International Conference on Artificial Intelligence and Statistics*, pp. 746-755, 2020.
9. C. Muge, T. Matthew, L. Ollie, A. E. Maraolo, J. Schafers and H. Antonia, “SARS-CoV-2, SARS-CoV, and MERS-CoV viral load dynamics, duration of viral shedding, and infectiousness: a systematic review and meta-analysis,” *The Lancet Microbe*, 2020.
10. S. Phipps, Q. Grafton and T. Kompas, “Robust estimates of the true (population) infection rate for COVID-19: a backcasting approach,” *Royal Society Open Science*, vol. 7, no. 11, 2020.
11. Ministry of Civil Protection and Ministry of Tourism, Hellenic Republic, “Protocol for Arrivals in Greece,” [Online]. Available: <https://travel.gov.gr/#/policy>.
12. R. Tibshirani, “Regression shrinkage and selection via the lasso,” *Journal of the Royal Statistical Society: Series B (Methodological)*, pp. 267-288, 1996.
13. S. Greenland and J. Robins, “Empirical-Bayes adjustments for multiple comparisons are sometimes useful,” *Epidemiology*, pp. 244-251, 1991.
14. O. J. Devine, T. Louis and E. Halloran, “Empirical Bayes methods for stabilizing incidence rates before mapping,” *Epidemiology*, pp. 622-630, 1994.
15. R. Dorfman, “The detection of defective members of large populations,” *The Annals of Mathematical Statistics*, vol. 14, no. 4, pp. 436-440, 1943.
16. W. Thompson, “On the likelihood that one unknown probability exceeds another in view of the evidence of two samples,” *Biometrika*, pp. 285-294, 1933.
17. T. L. Lai and H. Robbins, “Asymptotically efficient adaptive allocation rules,” *Advances in applied mathematics*, pp. 4-22, 1985.
18. J. Gittins, “Bandit processes and dynamic allocation indices,” *Journal of the Royal Statistical Society: Series B (Methodological)*, pp. 148-164, 1979.
19. P. Auer, “Using confidence bounds for exploitation-exploration trade-offs,” *Journal of Machine Learning Research*, pp. 397-422, 2002.
20. A. Tewari and S. A. Murphy, “From Ads to Interventions: Contextual Bandits in Mobile Health,” in *Mobile Health*, SpringerLink, 2017.
21. A. Durand, et al., “Contextual bandits for adapting treatment in a mouse model of de novo carcinogenesis,” in *Machine Learning for Healthcare Conference*, 2018.
22. L. Li, W. Chu, J. Langford and R. Schapire, “A contextual-bandit approach to personalized news article recommendation,” *Proceedings of the 19th international conference on World wide web*, pp. 6611-6670, 2010.
23. F. Amat, A. Chandrashekar, T. Jebara and J. Basilico, “Artwork personalization at Netflix,” *Proceedings of the 12th ACM conference on recommender systems*, pp. 487-488, 2018.
24. O. Besbes, Y. Gur and A. Zeevi, “Stochastic multi-armed-bandit problem with non-stationary rewards,” *Advances in neural information processing systems*, pp. 199-207, 2014.
25. H. Luo, C.-Y. Wei, A. Agarwal and J. Langford, “Efficient contextual bandits in non-stationary worlds,” *Conference on Learning Theory*, pp. 1739-1776, 2018.
26. H. Bastani and M. Bayati, “Online decision making with high-dimensional covariates,” *Operations Research*, pp. 276-294, 2020.
27. Z. Gao, Y. Han, Z. Ren and Z. Zhou, “Batched multi-armed bandits problem,” *Advances in Neural Information Processing Systems*, pp. 503-514, 2019.
28. V. Perchet, P. Rigollet, S. Chassang and E. Snowberg, “Batched bandit problems,” *The Annals of Statistics*, pp. 660-681, 2016.
29. S. Agrawal and N. Devanur, “Bandits with concave rewards and convex knapsacks,” *Proceedings of the fifteenth ACM conference on Economics and computation*, pp. 989-1006, 2014.
30. E. Gutin and V. Farias, “Optimistic gittins indices,” *Advances in Neural Information Processing Systems*, pp. 3153-3161, 2016.
31. Hellewell, J.; Abbott, S.; Gimma, A.; Bosse, N. I.; Jarvis, C. I.; Russell, T. W.; Munday, J. D.; Kucharski, A. J.; Edmunds, W. J.; Centre for the Mathematical Modelling of Infectious Diseases COVID-19 Working Group; Funk, S.; Eggo, R. M., “Feasibility of controlling COVID-19 outbreaks by isolation of cases and contacts,” *The Lancet. Global health*, 2020.
32. W. G. Imbens and B. D. Rubin, *Causal Inference in Statistics, Social and Biomedical Sciences*, Cambridge University Press, 2015.
33. P. Rosenbaum, and D. Rubin, “The central role of the propensity score in observational studies for causal effects,” *Biometrika*, vol. 70, no. 1, pp. 41-55, 1983.
34. Friedman, “Greedy function approximation: a gradient boosting machine,” *Annals of Statistics*, pp. 1189-1232, 2001.
35. N. Davies, et al., “Age-dependent effects in the transmission and control of COVID-19 epidemics,” *Nature Medicine*, 2020.
36. S. Davies and B. Bennet, “A gendered human rights analysis of Ebola and Zika: Locating gender in global health emergencies,” *International Affairs*.
37. S. Agrawal and N. Goyal, “Thompson sampling for contextual bandits with linear payoffs,” in *International Conference on Machine Learning*, 2013.
38. W. Chu, L. Li, L. Reyzin and R. Schapire, “Contextual bandits with linear payoff functions,” in *Conference on Artificial Intelligence and Statistics*, 2011.

**Publisher’s note** Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.


© The Author(s), under exclusive licence to Springer Nature Limited 2021




**Fig. 1 | Eva: A Reinforcement Learning System for COVID-19 Testing.** Arriving passengers submit travel and demographic information 24 hours prior to arrival. Based on these data and testing results from previous passengers, Eva selects a subset of passengers to test. Selected passengers

self-isolate for 24-48 hours while labs process samples. Positive passengers are then quarantined and contact tracing begins; negative passengers resume normal activities. Results are used to update Eva to improve future testing and maintain high-quality estimates of prevalence across traveler subpopulations.

ACCELERATED ARTICLE PREVIEW




**Fig. 2 | Comparing Eva vs. Randomized Surveillance Testing.** Infections caught by Eva (red) vs estimated number of cases caught by random, surveillance testing (teal). The peak (resp. off-peak) season is Aug. 6 to Oct. 1 (resp. Oct. 1 to Nov. 1) and is denoted with triangular (resp. circular) markers. Seasons are separated by the dotted line. Solid lines denote cubic-spline smoothing with 95% confidence intervals in grey.

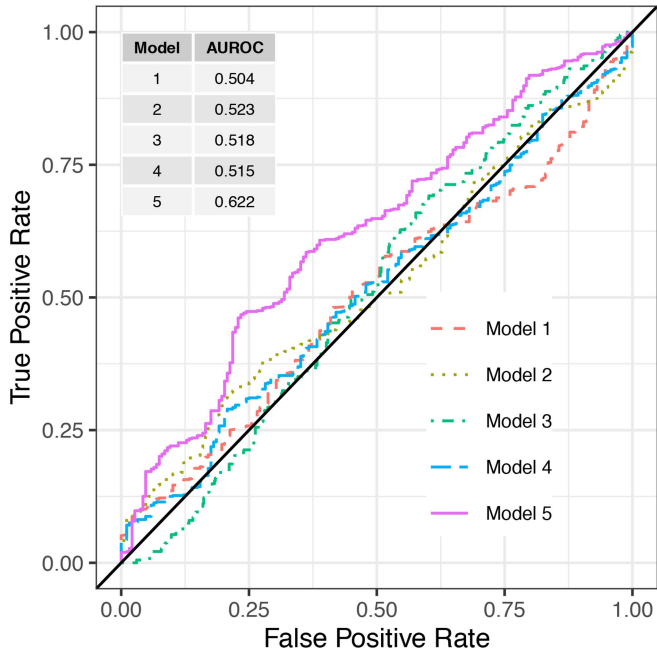


**Fig. 3 | Relative Efficacy of Eva over Random Surveillance vs. Testing Budget.** Ratio of number of infections caught by Eva relative to number of (estimated) infections caught by random surveillance testing, as a function of the fraction of tested travelers. Dotted (resp. dashed) line indicates the average fraction tested during the peak (resp. off-peak) tourist season. Triangular (circular) markers denote estimates from peak (off-peak) days. Solid blue line denotes cubic-spline smoothing with a 95% confidence interval in grey.





**Fig. 4 | Comparing Eva to Policies based on Epidemiological Metrics.** Lines represent cubic-spline smoothing of daily infections caught for each policy; raw points only shown for Eva and the “Cases” policy for clarity. The dotted line separates the peak (Aug. 6 to Oct. 1) and off-peak (Oct. 1 to Nov. 1) tourist seasons.



**Fig. 5 | Predictive Power of Publicly Reported Epidemiological Metrics.** Each of Models 1-4 uses a different subset of features from: 14-day time series of cases per capita, deaths per capita, tests performed per capita, and testing positivity rate. Model 5 additionally includes country fixed-effects to model country-level idiosyncratic behavior. Models 1-4 are essentially no better than random prediction, while Model 5 achieves slightly better performance. See Sec. 4.1 of Methods for details on model construction and features used in each model.

# Article

## Reporting summary

Further information on research design is available in the Nature Research Reporting Summary linked to this paper.

## Data availability

To support further research, aggregated, anonymized data are available at [https://github.com/kimondr/EVA\\_Public\\_Data](https://github.com/kimondr/EVA_Public_Data). These data aggregate passenger arrival and testing information over pairs of consecutive days, country of origin, and point of entry. The finer granularity data that support the (exact) findings of this study are protected by GDPR. These data are available from the Greek Ministry of Civil Protection but restrictions apply to the availability of these data, which were used under license for the current study, and so are not publicly available. Access to these data can only be granted by the Greek Ministry of Civil Protection ([info@gscp.gr](mailto:info@gscp.gr)) for research that is conducted in the public interest for public health (GDPR Recital 159) and scientific purposes (GDPR Article 89). Finally, the population-level epidemiological metrics used in our analysis can be obtained freely from the “Our World In Data COVID-19 dataset (<https://github.com/owid/covid-19-data/tree/master/public/data>).

## Code availability

All code used in this paper was written in a combination of R and Python 3.7. The code for the deployment of the algorithm on a sample

dataset is available at <https://github.com/vgupta1/EvaTargetedCovid-19Testing>. The code for reproducing the results of our counterfactual analysis is available at [https://github.com/vgupta1/Eva\\_CounterfactualAnalysis](https://github.com/vgupta1/Eva_CounterfactualAnalysis).

**Acknowledgements** The authors would like to thank all members of the Greek COVID-19 Taskforce, the Greek Prime Minister Kyriakos Mitsotakis, the Ministry of Digital Governance, the General Secretariat for Civil Protection, the Ministry of Health, the National Public Health Organization, the development team from Cytech as well as the border control agents, doctors, nurses and lab personnel that contributed to Eva’s deployment. Furthermore, the authors would like to thank Osbert Bastani for discussions and analysis on constructing custom risk metrics from public data. V.G. was partially supported by the National Science Foundation through NSF Grant CMMI-1661732. We also thank the Editor and 4 anonymous reviewers for their comments on an earlier draft of this manuscript.

**Author contributions** H.B., K.D., and V.G. constructed the model, designed and coded the algorithm, and performed the analysis in this paper. J.V. designed the software architecture and APIs to communicate with the Central Database of the Ministry of Digital Governance. C.H., P.L., G.M., D.P., and S.T. contributed to and informed epidemiological modeling choices of the system. All authors coordinated Eva’s operations and logistics throughout its deployment.

**Competing interests** H.B., V.G., and J.V. declare no conflict of interest. K.D. declares non-financial competing interest as an unpaid Data Science and Operations Advisor to the Greek Government from May 1<sup>st</sup>, 2020 to Nov 1<sup>st</sup>, 2020. C.H., P.L., G.M., D.P., and S.T. declare non-financial competing interest as members of the Greek National COVID-19 Taskforce.

## Additional information

**Supplementary information** The online version contains supplementary material available at <https://doi.org/10.1038/s41586-021-04014-z>.

**Correspondence and requests for materials** should be addressed to K.D.

**Peer review information** Nature thanks Mark Jit, Javier Sanchez-Monedero, Ziad Obermeyer and the other, anonymous, reviewer for their contribution to the peer review of this work.

**Reprints and permissions information** is available at <http://www.nature.com/reprints>.

ACCELERATED ARTICLE PREVIEW

## Reporting Summary

Nature Portfolio wishes to improve the reproducibility of the work that we publish. This form provides structure for consistency and transparency in reporting. For further information on Nature Portfolio policies, see our [Editorial Policies](#) and the [Editorial Policy Checklist](#).

### Statistics

For all statistical analyses, confirm that the following items are present in the figure legend, table legend, main text, or Methods section.

n/a Confirmed

- The exact sample size ( $n$ ) for each experimental group/condition, given as a discrete number and unit of measurement
- A statement on whether measurements were taken from distinct samples or whether the same sample was measured repeatedly
- The statistical test(s) used AND whether they are one- or two-sided  
*Only common tests should be described solely by name; describe more complex techniques in the Methods section.*
- A description of all covariates tested
- A description of any assumptions or corrections, such as tests of normality and adjustment for multiple comparisons
- A full description of the statistical parameters including central tendency (e.g. means) or other basic estimates (e.g. regression coefficient) AND variation (e.g. standard deviation) or associated estimates of uncertainty (e.g. confidence intervals)
- For null hypothesis testing, the test statistic (e.g.  $F$ ,  $t$ ,  $r$ ) with confidence intervals, effect sizes, degrees of freedom and  $P$  value noted  
*Give  $P$  values as exact values whenever suitable.*
- For Bayesian analysis, information on the choice of priors and Markov chain Monte Carlo settings
- For hierarchical and complex designs, identification of the appropriate level for tests and full reporting of outcomes
- Estimates of effect sizes (e.g. Cohen's  $d$ , Pearson's  $r$ ), indicating how they were calculated

*Our web collection on [statistics for biologists](#) contains articles on many of the points above.*

### Software and code

Policy information about [availability of computer code](#)

Data collection

Our system was deployed on Amazon Web Services (AWS). Code written in Python 3.7 and R version 4.0.5. The code for the algorithm and for the Counterfactual Analysis is posted on Github: <https://github.com/vgupta1/EvaTargetedCovid19Testing> and at [https://github.com/vgupta1/Eva\\_CounterfactualAnalysis](https://github.com/vgupta1/Eva_CounterfactualAnalysis), respectively.

Data analysis

Data Analysis was performed using Python 3.7 and R version 4.0.5

For manuscripts utilizing custom algorithms or software that are central to the research but not yet described in published literature, software must be made available to editors and reviewers. We strongly encourage code deposition in a community repository (e.g. GitHub). See the Nature Portfolio [guidelines for submitting code & software](#) for further information.

### Data

Policy information about [availability of data](#)

All manuscripts must include a [data availability statement](#). This statement should provide the following information, where applicable:

- Accession codes, unique identifiers, or web links for publicly available datasets
- A description of any restrictions on data availability
- For clinical datasets or third party data, please ensure that the statement adheres to our [policy](#)

To support further research, aggregated, anonymized data are available at [https://github.com/kimondr/EVA\\_Public\\_Data](https://github.com/kimondr/EVA_Public_Data). These data aggregate passenger arrival and testing information over pairs of consecutive days, country of origin, and point of entry.

The finer granularity data that support the (exact) findings of this study are protected by GDPR. These data are available from the Greek Ministry of Civil Protection but restrictions apply to the availability of these data, which were used under license for the current study, and so are not publicly available. Access to these data

can only be granted by the Greek Ministry of Civil Protection (info@gscp.gr) for research that is conducted in the public interest for public health (GDPR Recital 159) and scientific purposes (GDPR Article 89).

Finally, the population-level epidemiological metrics used in our analysis can be obtained freely from the "Our World In Data COVID-19 dataset (<https://github.com/owid/covid-19-data/tree/master/public/data>).

## Field-specific reporting

Please select the one below that is the best fit for your research. If you are not sure, read the appropriate sections before making your selection.

Life sciences       Behavioural & social sciences       Ecological, evolutionary & environmental sciences

For a reference copy of the document with all sections, see [nature.com/documents/nr-reporting-summary-flat.pdf](https://www.nature.com/documents/nr-reporting-summary-flat.pdf)

## Life sciences study design

All studies must disclose on these points even when the disclosure is negative.

Sample size	On average, 38,500 ( $\pm 13,590$ ) PLFs were processed each day (with higher traffic in August and September), and 18.3% ( $\pm 6.1\%$ ) of arriving households per day were tested (accounting for no-shows).  The effective sample size in our context was determined by the testing availability at each point of entry and was pre-specified by the General Secretariat of Public Health at the Greek Ministry of Health.  As evidenced by our confidence intervals, these sample sizes were indeed sufficient.
Data exclusions	0.6% of arrivals were tested by border agents at their own discretion (i.e., not recommended by Eva) and we drop these 0.6% of arrivals from the analysis. The exclusion criteria were pre-established.
Replication	We found that our algorithm consistently improved performance compared to benchmark policies throughout the summer. However, replication beyond the Greek border control context is out of the scope of our study.
Randomization	Samples were allocated based on our reinforcement learning algorithm. A detailed description is provided in Methods.
Blinding	Blinding is not relevant to this study. Border agents performed tests on individuals and (pseudonymized) samples were sent to laboratories that produced the result (positive or negative) for each traveler. The investigators (authors of this paper) were only provided pseudonymized, aggregate results.

## Reporting for specific materials, systems and methods

We require information from authors about some types of materials, experimental systems and methods used in many studies. Here, indicate whether each material, system or method listed is relevant to your study. If you are not sure if a list item applies to your research, read the appropriate section before selecting a response.

### Materials & experimental systems

n/a	Involvement in the study
<input checked="" type="checkbox"/>	<input type="checkbox"/> Antibodies
<input checked="" type="checkbox"/>	<input type="checkbox"/> Eukaryotic cell lines
<input checked="" type="checkbox"/>	<input type="checkbox"/> Palaeontology and archaeology
<input checked="" type="checkbox"/>	<input type="checkbox"/> Animals and other organisms
<input checked="" type="checkbox"/>	<input type="checkbox"/> Human research participants
<input checked="" type="checkbox"/>	<input type="checkbox"/> Clinical data
<input checked="" type="checkbox"/>	<input type="checkbox"/> Dual use research of concern

### Methods

n/a	Involvement in the study
<input checked="" type="checkbox"/>	<input type="checkbox"/> ChIP-seq
<input checked="" type="checkbox"/>	<input type="checkbox"/> Flow cytometry
<input checked="" type="checkbox"/>	<input type="checkbox"/> MRI-based neuroimaging