



Bayesian Best-Arm Identification for Selecting Influenza Mitigation Strategies

Pieter J. K. Libin^{1,2}(✉), Timothy Verstraeten¹, Diederik M. Roijers¹, Jelena Grujic¹, Kristof Theys², Philippe Lemey², and Ann Nowé¹

¹ Artificial Intelligence Lab, Department of Computer Science, Vrije Universiteit Brussel, Brussels, Belgium

{pieter.libin,timothy.verstraeten,jelena.grujic,ann.nowe}@vub.be

² Rega Institute for Medical Research, Clinical and Epidemiological Virology, KU Leuven - University of Leuven, Leuven, Belgium

{pieter.libin,kristof.theys,philippe.lemey}@kuleuven.be

Abstract. Pandemic influenza has the epidemic potential to kill millions of people. While various preventive measures exist (i.e., vaccination and school closures), deciding on strategies that lead to their most effective and efficient use remains challenging. To this end, individual-based epidemiological models are essential to assist decision makers in determining the best strategy to curb epidemic spread. However, individual-based models are computationally intensive and it is therefore pivotal to identify the optimal strategy using a minimal amount of model evaluations. Additionally, as epidemiological modeling experiments need to be planned, a computational budget needs to be specified a priori. Consequently, we present a new sampling technique to optimize the evaluation of preventive strategies using fixed budget best-arm identification algorithms. We use epidemiological modeling theory to derive knowledge about the reward distribution which we exploit using Bayesian best-arm identification algorithms (i.e., Top-two Thompson sampling and BayesGap). We evaluate these algorithms in a realistic experimental setting and demonstrate that it is possible to identify the optimal strategy using only a limited number of model evaluations, i.e., 2-to-3 times faster compared to the uniform sampling method, the predominant technique used for epidemiological decision making in the literature. Finally, we contribute and evaluate a statistic for Top-two Thompson sampling to inform the decision makers about the confidence of an arm recommendation. Code related to this paper is available at: <https://plibin-vub.github.io/epidemic-bandits>.

Keywords: Pandemic influenza · Multi-armed bandits
Fixed budget best-arm identification · Preventive strategies
Individual-based models

Electronic supplementary material The online version of this chapter (https://doi.org/10.1007/978-3-030-10997-4_28) contains supplementary material, which is available to authorized users.

1 Introduction

The influenza virus is responsible for the deaths of half of a million people each year. In addition, seasonal influenza epidemics cause a significant economic burden. While transmission is primarily local, a newly emerging variant may spread to pandemic proportions in a fully susceptible host population [29]. Pandemic influenza occurs less frequently than seasonal influenza but the outcome with respect to morbidity and mortality can be much more severe, potentially killing millions of people worldwide [29]. Therefore, it is essential to study mitigation strategies to control influenza pandemics.

For influenza, different preventive measures exist: i.e., vaccination, social measures (e.g., school closures and travel restrictions) and antiviral drugs. However, the efficiency of strategies greatly depends on the availability of preventive compounds, as well as on the characteristics of the targeted epidemic. Furthermore, governments typically have limited resources to implement such measures. Therefore, it remains challenging to formulate public health strategies that make effective and efficient use of these preventive measures within the existing resource constraints.

Epidemiological models (i.e., compartment models and individual-based models) are essential to study the effects of preventive measures *in silico* [2, 17]. While individual-based models are usually associated with a greater model complexity and computational cost than compartment models, they allow for a more accurate evaluation of preventive strategies [11]. To capitalize on these advantages and make it feasible to employ individual-based models, it is essential to use the available computational resources as efficiently as possible.

In the literature, a set of possible preventive strategies is typically evaluated by simulating each of the strategies an equal number of times [7, 13, 15]. However, this approach is inefficient to identify the optimal preventive strategy, as a large proportion of computational resources will be used to explore sub-optimal strategies. Furthermore, a consensus on the required number of model evaluations per strategy is currently lacking [34] and we show that this number depends on the *hardness* of the evaluation problem. Additionally, we recognize that epidemiological modeling experiments need to be planned and that a computational budget needs to be specified a priori. Therefore, we present a novel approach where we formulate the evaluation of preventive strategies as a *best-arm identification* problem using a *fixed budget* of model evaluations. In this work, the budget choice is left to the discretion of the decision maker, as would be the case for any uniform evaluation.

As running an individual-based model is computationally intensive (i.e., minutes to hours, depending on the complexity of the model), minimizing the number of required model evaluations reduces the total time required to evaluate a given set of preventive strategies. This renders the use of individual-based models attainable in studies where it would otherwise not be computationally feasible. Additionally, reducing the number of model evaluations will free up computational resources in studies that already use individual-based models, capacitating researchers to explore a larger set of model scenarios. This is important, as

considering a wider range of scenarios increases the confidence about the overall utility of preventive strategies [35].

In this paper, we contribute a novel technique to evaluate preventive strategies as a fixed budget best-arm identification problem. We employ epidemiological modeling theory to derive assumptions about the reward distribution and exploit this knowledge using Bayesian algorithms. This new technique enables decision makers to obtain recommendations in a reduced number of model evaluations. We evaluate the technique in an experimental setting, where we aim to find the best vaccine allocation strategy in a realistic simulation environment that models an influenza pandemic on a large social network. Finally, we contribute and evaluate a statistic to inform the decision makers about the confidence of a particular recommendation.

2 Background

2.1 Pandemic Influenza and Vaccine Production

The primary preventive strategy to mitigate seasonal influenza is to produce vaccine prior to the epidemic, anticipating the virus strains that are expected to circulate. This vaccine pool is used to inoculate the population before the start of the epidemic. While it is possible to stockpile vaccines to prepare for seasonal influenza, this is not the case for influenza pandemics, as the vaccine should be specifically tailored to the virus that is the source of the pandemic. Therefore, before an appropriate vaccine can be produced, the responsible virus needs to be identified. Hence, vaccines will be available only in limited supply at the beginning of the pandemic [33]. In addition, production problems can result in vaccine shortages [10]. When the number of vaccine doses is limited, it is imperative to identify an optimal vaccine allocation strategy [28].

2.2 Modeling Influenza

There is a long tradition of using individual-based models to study influenza epidemics [2, 15, 17], as they allow for a more accurate evaluation of preventive strategies. A state-of-the-art individual-based model that has been the driver for many high impact research efforts [2, 17, 18], is FluTE [6]. FluTE implements a contact model where the population is divided into communities of households [6]. The population is organized in a hierarchy of social mixing groups where the contact intensity is inversely proportional with the size of the group (e.g., closer contact between members of a household than between colleagues). Additionally, FluTE implements an individual disease progression model that associates different disease stages with different levels of infectiousness. FluTE supports the evaluation of preventive strategies through the simulation of therapeutic interventions (i.e., vaccines, antiviral compounds) and non-therapeutic interventions (i.e., school closure, case isolation, household quarantine).

2.3 Bandits and Best-Arm Identification

The *multi-armed bandit game* [1] involves a K -armed bandit (i.e., a slot machine with K levers), where each arm A_k returns a reward r_k when it is pulled (i.e., r_k represents a sample from A_k 's reward distribution). A common use of the bandit game is to pull a sequence of arms such that the cumulative regret is minimized [20]. To fulfill this goal, the player needs to carefully balance between exploitation and exploration.

In this paper, the objective is to recommend the best arm A^* (i.e., the arm with the highest average reward μ^*), after a fixed number of arm pulls. This is referred to as the fixed budget best-arm identification problem [1], an instance of the pure-exploration problem [4]. For a given budget T , the objective is to minimize the *simple regret* $\mu^* - \mu_J$, where μ_J is the average reward of the recommended arm A_J , at time T [5]. Simple regret is inversely proportional to the probability of recommending the correct arm A^* [24].

3 Related Work

As we established that a computational budget needs to be specified a priori, our problem setting matches the fixed budget best-arm identification setting. This differs from settings that attempt to identify the best arm with a predefined confidence: i.e., racing strategies [12], strategies that exploit the confidence bound of the arms' means [25] and more recently fixed confidence best-arm identification algorithms [16]. We selected Bayesian fixed budget best-arm identification algorithms, as we aim to incorporate prior knowledge about the arms' reward distributions and use the arms' posteriors to define a statistic to support policy makers with their decisions. We refer to [21, 24], for a broader overview of the state of the art with respect to (Bayesian) best-arm identification algorithms.

Best-arm identification algorithms have been used in a large set of application domains: i.a., evaluation of response surfaces, the initialization of hyperparameters and traffic congestion.

While other algorithms exist to rank or select bandit arms, e.g. [30], best-arm identification is best approached using adaptive sampling methods [23], as the ones we study in this paper.

In preliminary work, we explored the potential of multi-armed bandits to evaluate prevention strategies in a regret minimization setting, using default strategies (i.e., ϵ -greedy and UCB1). We presented this work at the 'Adaptive Learning Agents' workshop hosted by the AAMAS conference [26]. This setting is however inadequate to evaluate prevention strategies *in silico*, as minimizing cumulative regret is sub-optimal to identify the best arm. Additionally, in this workshop paper, the experiments considered a small and less realistic population, and only analyzed a limited range of R_0 values that is not representative for influenza pandemics.

4 Methods

We formulate the evaluation of preventive strategies as a multi-armed bandit game with the aim of identifying the best arm using a fixed budget of model evaluations. The presented method is generic with respect to the type of epidemic that is modeled (i.e., pathogen, contact network, preventive strategies). The method is evaluated in the context of pandemic influenza in the next section¹.

4.1 Evaluating Preventive Strategies with Bandits

A *stochastic epidemiological model* E is defined in terms of a model configuration $c \in \mathcal{C}$ and can be used to evaluate a preventive strategy $p \in \mathcal{P}$. The result of a model evaluation is referred to as the *model outcome* (e.g., prevalence, proportion of symptomatic individuals, morbidity, mortality, societal cost). Evaluating the model E thus results in a sample of the model's *outcome distribution*:

$$\text{outcome} \sim E(c, p), \text{ where } c \in \mathcal{C} \text{ and } p \in \mathcal{P} \quad (1)$$

Our objective is to find the optimal preventive strategy (i.e., the strategy that minimizes the expected outcome) from a set of alternative strategies $\{p_1, \dots, p_K\} \subset \mathcal{P}$ for a particular configuration $c_0 \in \mathcal{C}$ of a stochastic epidemiological model, where c_0 corresponds to the studied epidemic. To this end, we consider a multi-armed bandit with $K = |\{p_1, \dots, p_K\}|$ arms. Pulling arm p_k corresponds to evaluating p_k by running a simulation in the epidemiological model $E(c_0, p_k)$. The bandit thus has preventive strategies as arms with reward distributions corresponding to the outcome distribution of a stochastic epidemiological model $E(c_0, p_k)$. While the parameters of the reward distribution are known (i.e., the parameters of the epidemiological model), it is intractable to determine the optimal reward analytically. Hence, we must learn about the outcome distribution via interaction with the epidemiological model. In this work, we consider prevention strategies of equal financial cost, which is a realistic assumption, as governments typically operate within budget constraints.

4.2 Outcome Distribution

As previously defined, the reward distribution associated with a bandit's arm corresponds to the outcome distribution of the epidemiological model that is evaluated when pulling that arm. Therefore, we are able to specify prior knowledge about the reward distribution using epidemiological modeling theory.

It is well known that a disease outbreak has two possible outcomes: either it is able to spread beyond a local context and becomes a fully established epidemic or it fades out [32]. Most stochastic epidemiological models reflect this reality and hence its epidemic size distribution is bimodal [32]. When evaluating preventive strategies, the objective is to determine the preventive strategy that is most

¹ Code is available at <https://github.com/plibin-vub/bandits>.

suitable to mitigate an established epidemic. As in practice we can only observe and act on established epidemics, epidemics that faded out in simulation would bias this evaluation. Consequently, it is necessary to focus on the mode of the distribution that is associated with the established epidemic. Therefore we censor (i.e., discard) the epidemic sizes that correspond to the faded epidemic. The size distribution that remains (i.e., the one that corresponds with the established epidemic) is approximately Gaussian [3].

In this study, we consider a scaled epidemic size distribution, i.e., the proportion of symptomatic infections. Hence we can assume bimodality of the full size distribution and an approximately Gaussian size distribution of the established epidemic. We verified experimentally that these assumptions hold for all the reward distributions that we observed in our experiments (see Sect. 5).

To censor the size distribution, we use a threshold that represents the number of infectious individuals that are required to ensure an outbreak will only fade out with a low probability.

4.3 Epidemic Fade-Out Threshold

For heterogeneous host populations (i.e., a population with a significant variance among individual transmission rates, as is the case for influenza epidemics [9, 14]), the number of secondary infections can be accurately modeled using a negative binomial *offspring distribution* $NB(R_0, \gamma)$ [27], where R_0 is the basic reproductive number (i.e., the number of infections that is, by average, generated by one single infection) and γ is a dispersion parameter that specifies the extent of heterogeneity. The probability of epidemic extinction p_{ext} can be computed by solving $g(s) = s$, where $g(s)$ is the probability generating function (pgf) of the offspring distribution [27]. For an epidemic where individuals are targeted with preventive measures (e.g., vaccination), we obtain the following pgf

$$g(s) = \text{pop}_c + (1 - \text{pop}_c) \left(1 + \frac{R_0}{\gamma} (1 - s)\right)^{-\gamma} \quad (2)$$

where pop_c signifies the random proportion of controlled individuals [27]. From p_{ext} we can compute a threshold T_0 to limit the probability of extinction to a cutoff ℓ [19].

4.4 Best-Arm Identification with a Fixed Budget

Our objective is to identify the best preventive strategy (i.e., the strategy that minimizes the expected outcome) out of a set of preventive strategies, for a particular configuration $c_0 \in \mathcal{C}$ using a fixed budget T of model evaluations. To find the best prevention strategy, it suffices to focus on the mean of the outcome distribution, as it is approximately Gaussian with an unknown yet small variance [3], as we confirm in our experiments (see Fig. 1).

Successive Rejects was the first algorithm to solve the best-arm identification in a fixed budget setting [1]. For a K -armed bandit, Successive Rejects operates in $(K - 1)$ phases. At the end of each phase, the arm with the lowest average

reward is discarded. Thus, at the end of phase $(K - 1)$ only one arm survives, and this arm is recommended.

Successive Rejects serves as a useful baseline, however, it has no support to incorporate any prior knowledge. Bayesian best-arm identification algorithms on the other hand, are able to take into account such knowledge by defining an appropriate prior and posterior on the arms' reward distribution. As we will show, such prior knowledge can increase the best-arm identification accuracy. Additionally, at the time an arm is recommended, the posteriors contain valuable information that can be used to formulate a variety of statistics helpful to assist decision makers. We consider two state-of-the-art Bayesian algorithms: BayesGap [21] and Top-two Thompson sampling [31]. For Top-two Thompson sampling, we derive a statistic based on the posteriors to inform decision makers about the confidence of an arm recommendation: the probability of success.

As we established in the previous section, each arm of our bandit has a reward distribution that is approximately Gaussian with unknown mean and variance. For the purpose of genericity, we assume an uninformative Jeffreys prior $(\sigma_k)^{-3}$ on (μ_k, σ_k^2) , which leads to the following posterior on μ_k at the n_k^{th} pull [22]:

$$\sqrt{\frac{n_k^2}{S_{k,n_k}}}(\mu_k - \bar{x}_{k,n_k}) \mid \bar{x}_{k,n_k}, S_{k,n_k} \sim \mathcal{T}_{n_k} \tag{3}$$

where \bar{x}_{k,n_k} is the reward mean, S_{k,n_k} is the total sum of squares and \mathcal{T}_{n_k} is the standard student t-distribution with n_k degrees of freedom.

BayesGap is a gap-based Bayesian algorithm [21]. The algorithm requires that for each arm A_k , a high-probability upper bound $U_k(t)$ and lower bound $L_k(t)$ is defined on the posterior of μ_k at each time step t . Using these bounds, the gap quantity

$$B_k(t) = \max_{l \neq k} U_l(t) - L_k(t) \tag{4}$$

is defined for each arm A_k . $B_k(t)$ represents an upper bound on the simple regret (as defined in Sect. 2.3). At each step t of the algorithm, the arm $J(t)$ that minimizes the gap quantity $B_k(t)$ is compared to the arm $j(t)$ that maximizes the upper bound $U_k(t)$. From $J(t)$ and $j(t)$, the arm with the highest confidence diameter $U_k(t) - L_k(t)$ is pulled. The reward that results from this pull is observed and used to update A_k 's posterior. When the budget is consumed, the arm

$$J(\operatorname{argmin}_{t \leq T} B_{J(t)}(t)) \tag{5}$$

is recommended. This is the arm that minimizes the simple regret bound over all times $t \leq T$.

In order to use BayesGap to evaluate preventive strategies, we contribute problem-specific bounds. Given our posteriors (Eq. 3), we define

$$\begin{aligned} U_k(t) &= \hat{\mu}_k(t) + \beta \hat{\sigma}_k(t) \\ L_k(t) &= \hat{\mu}_k(t) - \beta \hat{\sigma}_k(t) \end{aligned} \tag{6}$$

where $\hat{\mu}_k(t)$ and $\hat{\sigma}_k(t)$ are the respective mean and standard deviation of the posterior of arm A_k at time step t , and β is the exploration coefficient.

The amount of exploration that is feasible given a particular bandit game, is proportional to the available budget, and inversely proportional to the game’s complexity [21]. This complexity can be modeled taking into account the game’s hardness [1] and the variance of the rewards. We use the hardness quantity defined in [21]:

$$H_\epsilon = \sum_k H_{k,\epsilon}^{-2} \tag{7}$$

with arm-dependent hardness

$$H_{k,\epsilon} = \max\left(\frac{1}{2}(\Delta_k + \epsilon), \epsilon\right), \text{ where } \Delta_k = \max_{l \neq k}(\mu_l) - \mu_k \tag{8}$$

Considering the budget T , hardness H_ϵ and a generalized reward variance σ_G^2 over all arms, we define

$$\beta = \sqrt{\frac{T - 3K}{4H_\epsilon \sigma_G^2}} \tag{9}$$

Theorem 1 in the Supplementary Information (Sect. 2) formally proves that using these bounds results in a probability of simple regret that asymptotically reaches the exponential lower bound of [21].

As both H_ϵ and σ_G^2 are unknown, in order to compute β , these quantities need to be estimated. Firstly, we estimate H_ϵ ’s upper bound \hat{H}_ϵ by estimating Δ_k as follows

$$\hat{\Delta}_k = \max_{1 \leq l < K; l \neq k} (\hat{\mu}_l(t) + 3\hat{\sigma}_l(t)) - (\hat{\mu}_k(t) - 3\hat{\sigma}_k(t)) \tag{10}$$

as in [21], where $\hat{\mu}_k(t)$ and $\hat{\sigma}_k(t)$ are the respective mean and standard deviation of the posterior of arm A_k at time step t . Secondly, for σ_G^2 we need a measure of variance that is representative for the reward distribution of all arms. To this end, when the arms are initialized, we observe their sample variance s_k^2 , and compute their average \bar{s}_G^2 .

As our bounds depend on the standard deviation $\hat{\sigma}_k(t)$ of the t-distributed posterior, each arm’s posterior needs to be initialized 3 times (i.e., by pulling the arm) to ensure that $\hat{\sigma}_k(t)$ is defined, this initialization also ensures proper posteriors [22].

Top-two Thompson sampling is a reformulation of the Thompson sampling algorithm, such that it can be used in a pure-exploration context [31]. Thompson sampling operates directly on the arms’ posterior of the reward distribution’s mean μ_k . At each time step, Thompson sampling obtains one sample for each arm’s posterior. The arm with the highest sample is pulled, and its reward is subsequently used to update that arm’s posterior. While this approach has been proven highly successful to minimize cumulative regret [8, 22], as it balances the exploration-exploitation trade-off, it is sub-optimal to identify the best arm [4]. To adapt Thompson sampling to minimize simple regret, Top-two Thompson sampling increases the amount of exploration. To this end, an exploration probability ω needs to be specified. At each time step, one sample is obtained for

each arm’s posterior. The arm A_{top} with the highest sample is only pulled with probability ω . With probability $1 - \omega$ we repeat sampling from the posteriors until we find an arm $A_{\text{top-2}}$ that has the highest posterior sample and where $A_{\text{top}} \neq A_{\text{top-2}}$. When the arm $A_{\text{top-2}}$ is found, it is pulled and the observed reward is used to update the posterior of the pulled arm. When the available budget is consumed, the arm with the highest average reward is recommended.

As Top-two Thompson sampling only requires samples from the arms’ posteriors, we can use the t-distributed posteriors from Eq. 3 as is. To avoid improper posteriors, each arm needs to be initialized 2 times [22].

As specified in the previous subsection, the reward distribution is censored. We observe each reward, but only consider it to update the arm’s value when it exceeds the threshold T_0 (i.e., when we receive a sample from the mode of the epidemic that represents the established epidemic).

4.5 Probability of Success

The probability that an arm recommendation is correct presents a useful confidence statistic to support policy makers with their decisions. As Top-two Thompson sampling recommends the arm with the highest average reward, and we assume that the arm’s reward distributions are independent, the probability of success is

$$P(\mu_J = \max_{1 \leq k \leq K} \mu_k) = \int_{x \in \mathbb{R}} \left[\prod_{k \neq J}^K F_{\mu_k}(x) \right] f_{\mu_J}(x) dx \quad (11)$$

where μ_J is the random variable that represents the mean of the recommended arm’s reward distribution, f_{μ_J} is the recommended arm’s posterior probability density function and F_{μ_k} is the other arms’ cumulative density function (full derivation in Supplementary Information, Sect. 3). As this integral cannot be computed analytically, we estimate it using Gaussian quadrature.

It is important to note that, while aiming for generality, we made some conservative assumptions: the reward distributions are approximated as Gaussian and the uninformative Jeffreys prior is used. These assumptions imply that the derived probability of success will be an under-estimator for the actual recommendation success, which is confirmed in our experiments.

5 Experiments

We composed and performed an experiment in the context of pandemic influenza, where we analyze the mitigation strategy to vaccinate a population when only a limited number of vaccine doses is available (details about the rationale behind this scenario in Sect. 2.1). In our experiments, we accommodate a realistic setting to evaluate vaccine allocation, where we consider a large and realistic social network and a wide range of R_0 values.

We consider the scenario when a pandemic is emerging in a particular geographical region and vaccines becomes available, albeit in a limited number of

doses. When the number of vaccine doses is limited, it is imperative to identify an optimal vaccine allocation strategy [28]. In our experiment, we explore the allocation of vaccines over five different age groups, that can be easily approached by health policy officials: pre-school children, school-age children, young adults, older adults and the elderly, as proposed in [6].

5.1 Influenza Model and Configuration

The epidemiological model used in the experiments is the FluTE stochastic individual-based model. In our experiment we consider the population of Seattle (United States) that includes 560,000 individuals [6]. This population is realistic both with respect to the number of individuals and its community structure, and provides an adequate setting for the validation of vaccine strategies [34] (more detail about the model choice in the Supplementary Information, Sect. 4).

At the first day of the simulated epidemic, 10 random individuals are seeded with an infection (more detail about the seeding choice in the Supplementary Information, Sect. 5). The epidemic is simulated for 180 days, during which no more infections are seeded. Thus, all new infections established during the run time of the simulation, result from the mixing between infectious and susceptible individuals. We assume no pre-existing immunity towards the circulating virus variant. We choose the number of vaccine doses to allocate to be approximately 4.5% of the population size [28].

We perform our experiment for a set of R_0 values within the range of 1.4 to 2.4, in steps of 0.2. This range is considered representative for the epidemic potential of influenza pandemics [2,28]. We refer to this set of R_0 values as \mathcal{R}_0 .

Note that the setting described in this subsection, in conjunction with a particular R_0 value, corresponds to a model configuration (i.e., $c_0 \in \mathcal{C}$).

The computational complexity of FluTE simulations depends both on the size of the susceptible population and the proportion of the population that becomes infected. For the population of Seattle, the simulation run time was up to $11\frac{1}{2}$ min (median of $10\frac{1}{2}$ min, standard deviation of 6s), on state-of-the-art hardware (details in Supplementary Information, Sect. 6).

5.2 Formulating Vaccine Allocation Strategies

We consider 5 age groups to which vaccine doses can be allocated: pre-school children (i.e., 0–4 years old), school-age children (i.e., 5–18 years old), young adults (i.e., 19–29 years old), older adults (i.e., 30–64 years old) and the elderly (i.e., >65 years old) [6]. An allocation scheme can be encoded as a Boolean 5-tuple, where each position in the tuple corresponds to the respective age group. The Boolean value at a particular position in the tuple denotes whether vaccines should be allocated to the respective age group. When vaccines are to be allocated to a particular age group, this is done proportional to the size of the population that is part of this age group [28]. To decide on the best vaccine allocation strategy, we enumerate all possible combinations of this tuple.

5.3 Outcome Distributions

To establish a proxy for the ground truth concerning the outcome distributions of the 32 considered preventive strategies, all strategies were evaluated 1000 times, for each of the R_0 values in \mathcal{R}_0 . We will use this ground truth as a reference to validate the correctness of the recommendations obtained throughout our experiments.

\mathcal{R}_0 presents us with an interesting evaluation problem. To demonstrate this, we visualize the outcome distribution for $R_0 = 1.4$ and for $R_0 = 2.4$ in Fig. 1 (the outcome distributions for the other R_0 values are shown in Sect. 7 of the Supplementary Information). Firstly, we observe that for different values of R_0 , the distances between top arms' means differ. Additionally, outcome distribution variances vary over the set of R_0 values in \mathcal{R}_0 . These differences produce distinct levels of evaluation hardness (see Sect. 4.4), and demonstrate the setting's usefulness as benchmark to evaluate preventive strategies. While we discuss the hardness of the experimental settings under consideration, it is important to state that our best-arm identification framework requires no prior knowledge on the problem's hardness. Secondly, we expect the outcome distribution to be bimodal. However, the probability to sample from the mode of the outcome distribution that represents the non-established epidemic decreases as R_0 increases [27]. This expectation is confirmed when we inspect Fig. 1, the left panel shows a bimodal distribution for $R_0 = 1.4$, while the right panel shows a unimodal outcome distribution for $R_0 = 2.4$, as only samples from the established epidemic were obtained.




Fig. 1. Violin plot that depicts the density of the outcome distribution (i.e., epidemic size) for 32 vaccine allocation strategies (left panel $R_0 = 1.4$, right panel $R_0 = 2.4$).

Our analysis identified that the best vaccine allocation strategy was $\langle 0, 1, 0, 0, 0 \rangle$ (i.e., allocate vaccine to school children, strategy 8) for all R_0 values in \mathcal{R}_0 .

5.4 Best-Arm Identification Experiment

To assess the performance of the different best-arm identification algorithms (i.e., Successive Rejects, BayesGap and Top-two Thompson sampling) we run each algorithm for all budgets in the range of 32 to 500. This evaluation is performed on the influenza bandit game that we defined earlier. For each budget, we run the algorithms 100 times, and report the recommendation success rate. In the previous section, the optimal vaccine allocation strategy was identified to be $\langle 0, 1, 0, 0, 0 \rangle$ for all R_0 in \mathcal{R}_0 . We thus consider a recommendation to be correct when it equals this vaccine allocation strategy.

We evaluate the algorithm’s performance with respect to each other and with respect to uniform sampling, the current state-of-the art to evaluate preventive strategies. The uniform sampling method pulls arm A_u for each step t of the given budget T , where A_u ’s index u is sampled from the uniform distribution $\mathcal{U}(1, K)$. To consider different levels of hardness, we perform this analysis for each R_0 value in \mathcal{R}_0 .

For the Bayesian best-arm identification algorithms, the prior specifications are detailed in Sect. 4.4. BayesGap requires an upper and lower bound that is defined in terms of the used posteriors. In our experiments, we use upper bound $U_k(t)$ and lower bound $L_k(t)$ that were established in Sect. 4.4. Top-two Thompson sampling requires a parameter that modulates the amount of exploration ω . As it is important for best-arm identification algorithms to differentiate between the top two arms, we choose $\omega = 0.5$, such that, in the limit, Top-two Thompson sampling will explore the top two arms uniformly.

We censor the reward distribution based on the threshold T_0 we defined in Sect. 4.3. This threshold depends on basic reproductive number R_0 and dispersion parameter γ . R_0 is defined explicitly for each of our experiments. For the dispersion parameter we choose $\gamma = 0.5$, which is a conservative choice according to the literature [9, 14]. We define the probability cutoff $\ell = 10^{-10}$.

Figure 2 shows recommendation success rate for each of the best-arm identification algorithms for $R_0 = 1.4$ (left panel) and $R_0 = 2.4$ (right panel). The results for the other R_0 values are visualized in Sect. 8 of the Supplementary Information. To complement these results, we show the recommendation success rate with confidence intervals in Sect. 9 of the Supplementary Information. The results for different values of R_0 clearly indicate that our selection of best-arm identification algorithms significantly outperforms the uniform sampling method. Overall, the uniform sampling method requires more than double the amount of evaluations to achieve a similar recommendation performance. For the harder problems (e.g., setting with $R_0 = 2.4$), recommendation uncertainty remains considerable even after consuming 3 times the budget required by Top-two Thompson sampling.

All best-arm identification algorithms require an initialization phase in order to output a well-defined recommendation. Successive Rejects needs to pull each arm at least once, while Top-two Thompson sampling and BayesGap need to pull each arm respectively 2 and 3 times (details in Sect. 4.4). For this reason, these algorithms’ performance can only be evaluated after this initialization

phase. BayesGap’s performance is on par with Successive Rejects, except for the hardest setting we studied (i.e., $R_0 = 2.4$). In comparison, Top-two Thompson sampling consistently outperforms Successive Rejects 30 pulls after the initialization phase. Top-two Thompson sampling needs to initialize each arm’s posterior with 2 pulls, i.e., double the amount of uniform sampling and Successive Rejects. However, our experiments clearly show that none of the other algorithms reach any acceptable recommendation rate using less than 64 pulls.




Fig. 2. In this figure, we present the results for the experiment with $R_0 = 1.4$ (left panel) and $R_0 = 2.4$ (right panel). Each curve represents the rate of successful arm recommendations (y-axis) for a range of budgets (x-axis). A curve is shown for each of the considered algorithms: BayesGap (legend: BG), Successive Rejects (legend: SR), Top-two Thompson sampling (legend: TtTs) and uniform sampling (legend: Uni).

In Sect. 4 we derived a statistic to express the probability of success (P_s) concerning a recommendation made by Top-two Thompson sampling. We analyzed this probability for all the Top-two Thompson sampling recommendations that were obtained in the experiment described above. To provide some insights on how this statistic can be used to support policy makers, we show the P_s values of all Top-two Thompson sampling recommendations for $R_0 = 2.4$ in the left panel of Fig. 3 (Figures for the other R_0 values in Sect. 10 of the Supplementary Information). This Figure indicates that P_s closely follows recommendation correctness and that the uncertainty of P_s is inversely proportional to the size of the available budget. Additionally, in the right panel of Fig. 3 (Figures for the other R_0 values in Sect. 11 of the Supplementary Information) we confirm that P_s underestimates recommendation correctness. These observations show that P_s has the potential to serve as a conservative statistic to inform policy makers about the confidence of a particular recommendation, and thus can be used to define meaningful cutoffs to guide policy makers in their interpretation of the recommendation of preventive strategies.

In this work, we define uninformed priors to ensure a generic framework. This does not exclude decision makers to use priors that include more domain knowledge (e.g., dependence between arms), if this is available. We do however show in our experiments that the use of these uninformed priors lead to a significant performance increase.




Fig. 3. Top-two Thompson sampling was run 100 times for each budget for the experiment with $R_0 = 2.4$. For each of the recommendations, P_s was computed. In the left panel, these P_s values are shown as a scatter plot, where each point’s color reflects the correctness of the recommendation (see legend). In the right panel, the P_s values were binned (i.e., 0.5 to 1 in steps of 0.05). Per bin, we thus have a set of Bernoulli trials, for which we show the empirical success rate (blue scatter) and the Clopper-Pearson confidence interval (blue confidence bounds). The orange reference line denotes perfect correlation between the empirical success rate and the estimated probability of success. (Color figure online)

6 Conclusion

We formulate the objective to select the best preventive strategy in an individual-based model as a fixed budget best-arm identification problem. We set up an experiment to evaluate this setting in the context of a realistic influenza pandemic. To assess the best arm recommendation performance of the preventive bandit, we report a success rate over 100 independent bandit runs.

We demonstrate that it is possible to efficiently identify the optimal preventive strategy using only a limited number of model evaluations, even if there is a large number of preventive strategies to consider. Compared to uniform sampling, our technique is able to recommend the best preventive strategy reducing the number of required model evaluations 2-to-3 times, when using Top-two Thompson sampling. Additionally, we defined a statistic to support policy makers with their decisions, based on the posterior information obtained during Top-two Thompson sampling. As such, we present a decision support tool to assist policy makers to mitigate epidemics. Our framework will enable the use of individual-based models in studies where it would otherwise be computationally too prohibitive, and allow researchers to explore a wider variety of model scenarios.

Acknowledgments. Pieter Libin and Timothy Verstraeten were supported by a PhD grant of the FWO (Fonds Wetenschappelijk Onderzoek - Vlaanderen). Kristof Theys, Jelena Grujic and Diederik Roijers were supported by a postdoctoral grant of the FWO. The computational resources were provided by an EWI-FWO grant (Theys, KAN2012 1.5.249.12.). We thank the anonymous reviewers for their insightful comments that allowed us to improve this work.

References

1. Audibert, J.Y., Bubeck, S.: Best arm identification in multi-armed bandits. In: COLT-23th Conference on Learning Theory (2010)
2. Basta, N.E., Chao, D.L., Halloran, M.E., Matrajt, L., Longini, I.M.: Strategies for pandemic and seasonal influenza vaccination of schoolchildren in the United States. *Am. J. Epidemiol.* **170**(6), 679–686 (2009)
3. Britton, T.: Stochastic epidemic models: a survey. *Math. Biosci.* **225**(1), 24–35 (2010)
4. Bubeck, S., Munos, R., Stoltz, G.: Pure exploration in multi-armed bandits problems. In: Gavaldà, R., Lugosi, G., Zeugmann, T., Zilles, S. (eds.) ALT 2009. LNCS (LNAI), vol. 5809, pp. 23–37. Springer, Heidelberg (2009). https://doi.org/10.1007/978-3-642-04414-4_7
5. Bubeck, S., Munos, R., Stoltz, G.: Pure exploration in finitely-armed and continuous-armed bandits. *Theor. Comput. Sci.* **412**(19), 1832–1852 (2011)
6. Chao, D.L., Halloran, M.E., Obenchain, V.J., Longini Jr., I.M.: FluTE, a publicly available stochastic influenza epidemic simulation model. *PLoS Comput. Biol.* **6**(1), e1000656 (2010)
7. Chao, D.L., Halstead, S.B., Halloran, M.E., Longini, I.M.: Controlling Dengue with Vaccines in Thailand. *PLoS Negl. Trop. Dis.* **6**(10), e1876 (2012)
8. Chapelle, O., Li, L.: An empirical evaluation of Thompson sampling. In: *Advances in Neural Information Processing Systems*, pp. 2249–2257 (2011)
9. Dorigatti, I., Cauchemez, S., Pugliese, A., Ferguson, N.M.: A new approach to characterising infectious disease transmission dynamics from sentinel surveillance: application to the Italian 2009/2010 A/H1N1 influenza pandemic. *Epidemics* **4**(1), 9–21 (2012)
10. Enserink, M.: Crisis underscores fragility of vaccine production system. *Science* **306**(5695), 385 (2004)
11. Eubank, S., Kumar, V., Marathe, M., Srinivasan, A., Wang, N.: Structure of social contact networks and their impact on epidemics. *DIMACS Ser. Discrete Math. Theor. Comput. Sci.* **70**(0208005), 181 (2006)
12. Even-Dar, E., Mannor, S., Mansour, Y.: Action elimination and stopping conditions for the multi-armed bandit and reinforcement learning problems. *J. Mach. Learn. Res.* **7**(Jun), 1079–1105 (2006)
13. Ferguson, N.M., Cummings, D.A.T., Cauchemez, S., Fraser, C.: Others: strategies for containing an emerging influenza pandemic in Southeast Asia. *Nature* **437**(7056), 209 (2005)
14. Fraser, C., Cummings, D.A.T., Klinkenberg, D., Burke, D.S., Ferguson, N.M.: Influenza transmission in households during the 1918 pandemic. *Am. J. Epidemiol.* **174**(5), 505–514 (2011)
15. Fumanelli, L., Ajelli, M., Merler, S., Ferguson, N.M., Cauchemez, S.: Model-based comprehensive analysis of school closure policies for mitigating influenza epidemics and pandemics. *PLoS Comput. Biol.* **12**(1), e1004681 (2016)
16. Garivier, A., Kaufmann, E.: Optimal best arm identification with fixed confidence. In: *Conference on Learning Theory*, pp. 998–1027 (2016)
17. Germann, T.C., Kadau, K., Longini, I.M., Macken, C.A.: Mitigation strategies for pandemic influenza in the United States. *Proc. Nat. Acad. Sci. U.S.A.* **103**(15), 5935–5940 (2006)
18. Halloran, M.E., Longini, I.M., Nizam, A., Yang, Y.: Containing bioterrorist smallpox. *Science (New York, N.Y.)* **298**(5597), 1428–1432 (2002)

19. Hartfield, M., Alizon, S.: Introducing the outbreak threshold in epidemiology. *PLoS Pathog* **9**(6), e1003277 (2013)
20. Herbert, R.: Some aspects of the sequential design of experiments. *Bull. Am. Math. Soc.* **58**(5), 527–535 (1952)
21. Hoffman, M., Shahriari, B., Freitas, N.: On correlation and budget constraints in model-based bandit optimization with application to automatic machine learning. In: *Artificial Intelligence and Statistics*, pp. 365–374 (2014)
22. Honda, J., Takemura, A.: Optimality of Thompson sampling for Gaussian bandits depends on priors. In: *AISTATS*, pp. 375–383 (2014)
23. Jennison, C., Johnstone, I.M., Turnbull, B.W.: Asymptotically optimal procedures for sequential adaptive selection of the best of several normal means. *Stat. Decis. Theory Relat. Top.* **III 2**, 55–86 (1982)
24. Kaufmann, E., Cappé, O., Garivier, A.: On the complexity of best arm identification in multi-armed bandit models. *J. Mach. Learn. Res.* **17**(1), 1–42 (2016)
25. Kaufmann, E., Kalyanakrishnan, S.: Information complexity in bandit subset selection. In: *Conference on Learning Theory*, pp. 228–251 (2013)
26. Libin, P., Verstraeten, T., Theys, K., Roijsers, D.M., Vranx, P., Nowé, A.: Efficient evaluation of influenza mitigation strategies using preventive bandits. In: Sukthankar, G., Rodriguez-Aguilar, J.A. (eds.) *AAMAS 2017. LNCS (LNAI)*, vol. 10643, pp. 67–85. Springer, Cham (2017). https://doi.org/10.1007/978-3-319-71679-4_5
27. Lloyd-Smith, J.O., Schreiber, S.J., Kopp, P.E., Getz, W.M.: Superspreading and the effect of individual variation on disease emergence. *Nature* **438**(7066), 355–359 (2005)
28. Medlock, J., Galvani, A.P.: Optimizing influenza vaccine distribution. *Science* **325**(5948), 1705–1708 (2009)
29. Paules, C., Subbarao, K.: Influenza. *The Lancet* (2017)
30. Powell, W.B., Ryzhov, I.O.: *Optimal Learning*, vol. 841. Wiley, Hoboken (2012)
31. Russo, D.: Simple Bayesian algorithms for best arm identification. In: *Conference on Learning Theory*, pp. 1417–1418 (2016)
32. Watts, D.J., Muhamad, R., Medina, D.C., Dodds, P.S.: Multiscale, resurgent epidemics in a hierarchical metapopulation model. *Proc. Nat. Acad. Sci. U.S.A.* **102**(32), 11157–11162 (2005)
33. WHO: WHO guidelines on the use of vaccines and antivirals during influenza pandemics (2004)
34. Willem, L., Stijven, S., Vladislavleva, E., Broeckhove, J., Beutels, P., Hens, N.: Active learning to understand infectious disease models and improve policy making. *PLoS Comput. Biol.* **10**(4), e1003563 (2014)
35. Wu, J.T., Riley, S., Fraser, C., Leung, G.M.: Reducing the impact of the next influenza pandemic using household-based public health interventions. *PLoS Med.* **3**(9), e361 (2006)