Containing the Spread of a Contagion on a Tree

Michela Meister
Cornell University
Ithaca, NY, USA
meister@cs.cornell.edu

Jon Kleinberg
Cornell University
Ithaca, NY, USA
kleinberg@cornell.edu

ABSTRACT

Contact tracing can be thought of as a race between two processes: an infection process and a tracing process. In this paper, we study a simple model of infection spreading on a tree, and a tracer who stabilizes one node at a time. We focus on the question, how should the tracer choose nodes to stabilize so as to prevent the infection from spreading further? We study simple policies, which prioritize nodes based on time, infectiousness, or probability of generating new contacts.

KEYWORDS

contagion dynamics; contact tracing; prioritization policies; cascade processes

ACM Reference Format:
Michela Meister and Jon Kleinberg. 2024. Containing the Spread of a Contagion on a Tree. In *Proc. of the 23rd International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2024), Auckland, New Zealand, May 6 – 10, 2024*, IFAAMAS, 9 pages.

1 INTRODUCTION

Mathematical models have played an important role in epidemiology, providing tools and frameworks complementing empirical and public health research. One key example are branching processes, which lead to the development of the R_0 metric for measuring the spread of disease [19]. While there are many mathematical models of the spread of disease, far fewer models exist for contact tracing. Here we present an initial mathematical model of contact tracing which we use to explore algorithmic questions in designing contact tracing interventions.

Contact tracing is the iterative process of identifying individuals (the *contacts*) exposed to an infected case [3, 25, 31]. These contacts may then be tested for infection, treated, or quarantined, depending on the nature of the disease, to limit the spread of further infections.

This can be thought of as a race between two processes: an infection process and a tracing process. The goal of the tracing process is to identify infected cases faster than the disease spreads, so that eventually no new infections occur, ie the infection is *contained*. Contact tracing is often implemented by teams of human tracers, and as a result, the tracing process is limited by the number of human tracers available. A key strategic decision is how to maximize the effectiveness of this limited tracing capacity [16, 23, 29]. To simplify things, we model the tracing process as a single *tracer* who is given a list of contacts exposed to infection. We focus on the question, given a list of contacts exposed to infection, which contact should the tracer investigate or *query* next? In particular, how does the tracer's policy for querying contacts affect the probability that the infection is contained?

One of the challenges in studying these questions is the lack of simple models that manage to articulate trade-offs between the infection and tracing processes. In the contact tracing literature, there are few models which simultaneously capture the dynamics of an infection process and a resource-constrained tracing process. Recent surveys on the contact tracing literature specifically note that "few models take the limited capacity of the public health system into account" [23] or "consider...the practical constraints that resources for contact tracing and follow-up control measures might not be available at full throttle" [16]. Meanwhile, the literature on probabilistic models provides many epidemic models on trees, but these models do not consider the effect of a tracing process. Thus it seems as if a model describing the interaction between an infection process and a tracing process has been absent from these two fields. This provides an opportunity to apply AI to automate decision-making in a complex environment involving multiple agents.

Related work. A few other papers analyze contact tracing under resource constraints, however in somewhat different settings. In [20], Meister and Kleinberg develop a model of contact tracing in which the infection and tracing processes operate in two disjoint phases. In the first phase the infection spreads throughout the population; in the second phase the population is in "lockdown" and no new infections occur. Tracing proceeds in the second phase, and the tracer's objective is to identify infected nodes efficiently so as to maximize a total "benefit". In contrast, this paper studies concurrent infection and tracing processes, where the tracer's objective is to contain the spread of the infection.

Armbruster and Brandeau also study contact tracing under resource constraints in [1, 2]. In their model there are fixed resources to allocate across the two interventions, contact tracing and surveillance testing. The primary goal is to find the optimal allocation of resources so as to provide the best health outcomes for the population. They evaluate a few simple policies for prioritizing contacts in [2], and choose the policy that results in the lowest prevalence of infection for their main analysis in [1]. However, their main focus is on determining the best allocation of resources across these two systems. In comparison, our work focuses on analyzing and measuring the performance of different prioritization policies across a wide range of infection parameters. Li and collaborators study contact tracing when there is a limit on the number of individuals that can be quarantined in [17]. In [28], Guni and collaborators study the problem of inferring which individuals in a populatin are infected based on testing results for a subset of individuals and a contact graph. Finally, in a somewhat different context, Ben-Eliezer, Mossel,

Proc. of the 23rd International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2024), N. Alechina, V. Dignum, M. Dastani, J.S. Sichman (eds.), May 6 – 10, 2024, Auckland, New Zealand. © 2024 International Foundation for Autonomous Agents and Multiagent Systems (www.ifaamas.org).

and Sudan study a mathematical model of information spreading on a network with errors in communication and investigate approaches for error correction [4].

The model. To address our questions, we need to be able to define trade-offs between the tracer's policy for querying individuals and factors such as an individual's rate of meeting new contacts and the probability that they transmit the infection to a contact. To do this, we develop a simple model of contact tracing on a tree, which involves concurrent infection and tracing processes.

First we describe the infection process uninhibited by any tracing. Each individual is represented as a node with a binary infection status. A node v is governed by two parameters: the probability q_v that it meets a new contact and the probability p_v that it transmits the infection to a contact. These parameters are sampled independently for each node, with $p_v \sim D_p$ and $q_v \sim D_q$. We will discuss more about these distributions later on, but the problem is still interesting even when both distributions take just a single fixed value. Initially all nodes are uninfected. In round $t = 0$ a node r becomes infected with a probability drawn from D_p. In each round $t > 0$, each node v meets a new contact u with probability q_v. If v is infected, it infects u with probability p_v. This process generates a tree, where r is the root, the nodes in the first layer are r's contacts, the nodes in the second layer are contacts of those contacts, and so on. If a node v joins the tree in round t, its *time-of-arrival* is $\tau_v = t$. If root r is not infected at $t = 0$, the infection is trivially contained.

In order to define the tracing process, we assign each node a second binary status, indicating whether it is *active* or *stable*. A node that is *active* probabilistically generates new contacts at each round, as defined by the infection process. A node that is *stable* no longer generates new contacts and therefore cannot further spread the infection. Initially, every node is *active*.

Contact tracing starts once the infection is already underway, at time $t = k$, when the tracer identifies a root r as an index case. From then on the tracer selects one node to *query* at each step. Note that, while the tracer only queries one node at each step, we can change the rate of tracing relative to the infection process by changing the contact probability q. Increasing q causes the infection to spread more quickly thereby decreasing the relative rate of tracing, and decreasing q causes the infection to spread more slowly, thereby increasing the relative rate of tracing. Querying a node reveals its infection status, and if a node is infected, two events occur: (1) the node is stabilized and (2) the node's children (ie contacts) are revealed. Thus querying an infected node has two benefits: it prevents further infections and reveals individuals exposed to infection.[1] At step $t = k$ the only node the tracer may query is the root r. From then on the tracer may only query a node if its parent is an infected node queried on an earlier step.

We now describe the concurrent infection and tracing processes. We say that an *instance* of the contact tracing problem is defined by the three parameters D_p, D_q, and k. The process begins at step $t = 0$ when the root r becomes infected with a probability drawn from D_p. The infection process runs uninhibited for steps $0 \le t < k$. During each step $t \ge k$, first the tracer queries a node and then a single round of the infection process runs. During the infection round only active nodes generate new contacts. The infection is *contained* if the tracer stabilizes all infected nodes.

Policies for querying nodes. We can think of the tracing process as maintaining a subtree where each node in the subtree has an infected parent. The *frontier* is the set of all leaves in the subtree which have not yet been queried. We assume that the tracer observes the triple (p_u, q_u, τ_u) for each node in the subtree and that, for the purposes of querying, any two nodes in the frontier with the same triple of parameters are indistinguishable. A *policy* is any rule that dictates which node from the frontier is queried next. Note that if the frontier is empty, then all infected nodes have been stabilized, and therefore the infection is contained.

We say that a policy is *non-trivial* if it only queries the children of infected nodes. (Since the children of uninfected nodes are guaranteed to be uninfected, there is no reason to query them.) The remainder of the paper considers only non-trivial policies.

1.1 Summary and Overview of Results

We analyze the effectiveness of different tracing policies, with a primary focus on the following question.

Question 1. *How does the tracer's policy for querying nodes affect the probability that the infection is contained?*

We study this question via both theoretical analysis and computational experiments. To begin, we establish basic theoretical bounds in section 2, which characterize the performance of any non-trivial policy under certain conditions. In particular, we show that if either the infection probability p_v or the contact probability q_v is sufficiently small for all nodes, then any non-trivial policy contains the infection with high probability. On the other hand, we show that if both contagion parameters are sufficiently large, then every policy fails with high probability.

Thus the results in this first section focus on settings in which policy choice is inconsequential; either any non-trivial policy is likely to contain the infection, or no non-trivial policy is likely to contain the infection. This motivates the question of whether there exists an instance in which the choice of policy *is* significant.

Question 2. *Is there an instance and a pair of policies $\mathcal{A}_1$ and $\mathcal{A}_2$ so that the probability of containment under $\mathcal{A}_1$ is greater than the probability of containment under $\mathcal{A}_2$?*

This question is the primary focus of the remainder of the paper. To begin, we start with the simplest setting possible, where D_p and D_q are point mass distributions, that is, where all nodes have the same values of p and q. In such a setting, one might think that the probability of containment ought to be agnostic to the policy chosen by the tracer. However, two nodes with the same infection and contact parameters p and q may still differ in their time-of-arrival τ.

There are two natural policies for ordering nodes by time-of-arrival. The *ascending-time* policy orders nodes by ascending time-of-arrival, and the *descending-time* policy orders nodes by descending time-of-arrival. In section 3 we prove that for a specific choice of p and q, descending-time has a strictly higher probability of containment than ascending-time. Given this result, one might wonder whether descending-time is the better strategy in all instances. In section 4 we compare the performance of ascending-time

[1] If an individual is found to be uninfected, they do not need to be stabilized, since they cannot infect any contacts.

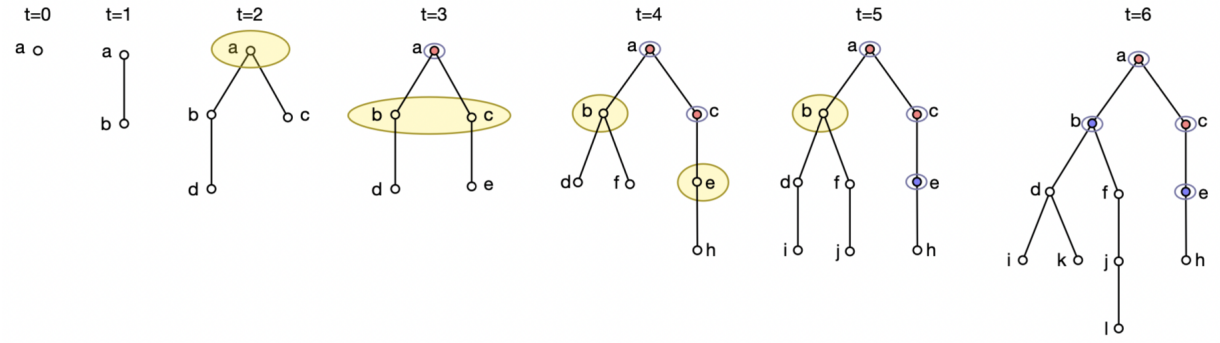

Figure 1: This example illustrates concurrent infection and tracing processes. The infection process begins at $t = 0$ and runs uninhibited for three steps. Tracing begins at $t = 3$. At the start of step $t = 3$, a is the only node in the frontier. The tracer queries a, and since a is infected, its children b and c join the frontier. Then another round of the infection process runs, in which every node that has not yet been queried probabilistically generates a new contact. In this case, c generates child e. At $t = 4$, the tracer queries c, which is infected, so e joins the frontier. Another round of the infection process runs, where b generates child f and e generates child h. At $t = 6$, the tracer queries b, the sole node in the frontier. Node b is uninfected, so the frontier is empty, and thus the infection is contained.

and descending-time for a large range of contagion parameters via computational experiments and find numerous instances in which we observe a significantly higher probability of containment for ascending-time. From these computational experiments, we find that each of the two policies commands a region of the parameter space of non-trivial size where it enjoys signficantly better performance, and we find that descending-time has the larger region. In section 5 we explore instances in which different nodes may have different contagion parameters. For these settings we study policies that prioritize nodes by descreasing infection probability p_v or by decreasing contact probability q_v.

Paper organization. We begin with a summary of related work. Section 2 establishes basic theoretical bounds. Section 3 shows a specific instance where policy choice provably affects probability of containment. Section 4 studies this question further, by comparing the performance of the ascending-time and descending-time policies via computational experiments. Section 5 compares the performance of other policies for prioritizing nodes beyond time-based methods. Finally, section 6 presents future work and open questions.

Further Related Work. Tian et al. study Tuberculosis contact tracing on a simulated network based on the population of Saskatchewan, Canada, and compare different prioritization policies for tracing individuals, with a particular focus on prioritizations based on patient demographics [30]. Prior work by Fraser et al. and Klinkenberg studies when an outbreak of a disease may be contained by tracing and isolation interventions, with a focus on HIV, smallpox, and influenza, among other diseases [7, 14]. Hellewell et al. study this question for COVID-19 specifically [8]. Kretzschmar et al. study the effect of time delays on contact tracing for COVID-19 via computational simulations [15]. Kwok et al. review models of contact tracing and call for more models to account for resource constraints in tracing [16]. Kaplan et al. model a tracing and vaccination response to a bioterrorism attack in [12, 13].

Muller et al. study contact tracing as a branching process [24]. Eames et al. study different contact tracing strategies for a compartmental model of infection [5, 9]. Eames and Keeling study the relationship between the fraction of contacts which are traced and the rate at which the disease spreads in the context of sexually transmitted diseases [6].

2 BASIC THEORETICAL BOUNDS

Recall that a non-trivial policy is one which only queries the children of infected nodes. Our basic theoretical bounds define conditions under which any non-trivial policy succeeds and under which any non-trivial policy fails. We focus on the following question.

Question 3. *Fix a non-trivial policy P. Under what conditions, with high probability, does P contain the infection? Under what conditions, with high probability, does P fail to contain the infection?*

Figure 2 outlines our results in this section. First we establish that, if either the infection probability p_v or contact probability q_v is sufficiently small for all nodes, then any non-trivial policy contains the infection with high probability, which is shown in theorems 1 and 2. On the other hand, if both the infection probability and contact probability are sufficiently large for all nodes, theorem 3 shows that for any fixed non-trivial policy P, with high probablity P does not contain the infection.

2.1 Conditions Under Which Containment is Likely

We present two theorems, which together show that if either the infection probability or contact probability is below a certain threshold for all nodes, then any non-trivial policy contains the infection with high probability.

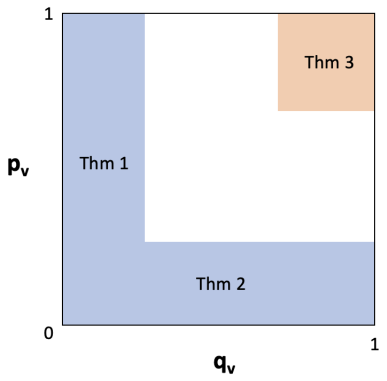

Figure 2: Our basic theoretical bounds describe parameter regimes in which the choice of policy is inconsequential. If the infection or contact probability is very low for all nodes, then any non-trivial policy contains the infection with high probability, as shown in theorems 1 and 2. However, if the infection and contact probabilities are both very high for all nodes, then containment is highly unlikely, regardless of the non-trivial policy chosen, as shown in theorem 3.

Defining a transcript. For the following two theorems it will be helpful to analyze the tracing process through the lens of deferred decisions. This analysis changes nothing about how the contact tracing process is defined, but simply makes it easier for us to analyze. First we will generate a *transcript* $T'_0, T'_1, \dots$ of a tree with the given contagion parameters growing uninhibited by any tracing. During the tracing process, we construct infection tree $T_0, T_1, \dots$ by "replaying" the transcript. For example, when a new round of infection occurs at time t, we refer to T'_t to determine the nodes to add to T_t and their infection statuses. The benefit of this framework is that we can prove claims about the transcript $T'_0, T'_1, \dots$, which is often much easier to analyze, and show that these claims hold for the infection tree $T_0, T_1, \dots$ as well.

To start, we show that if the contact probability q_v is sufficiently small for all nodes, then any non-trivial policy contains the infection with high probability.

Theorem 1. *Fix a failure probability $\delta \in (0,1)$ and an arrival time $k \in \mathbb{N}$. Suppose that each node v has infection probability $p_v \leq 1$. There is a $q(\delta,k) \in (0,1]$ such that, if each node v has contact probability $q_v < q(\delta,k)$, then any non-trival policy contains the infection with probability at least $1-\delta$.*

Proof Sketch. Recall that in each step, either the frontier is empty and therefore the infection is contained, or the tracer stabilizes one node from the frontier. Therefore, at any step $t > k$, either the infection is already contained, or the tracer has stabilized at least $t-k$ nodes.

Let $T'_0, T'_1, \dots$ be the transcript of a tree with the given contagion parameters. The proof idea is to show that there is a time $m > k$ at which point with probability at least $1-\delta$, T'_m has at most $m-k$ nodes. By definition, T_m has at most the number of nodes in T'_m. Therefore, with probability at least $1-\delta$, the tracer has contained the infection by time $t = m$. We use Markov's Inequality to bound the number of nodes in T'_m. **??** provides the full proof of theorem 1. □

Similarly, the following theorem demonstrates that if for all nodes the probability of infection is below a certain threshold, then any non-trivial policy contains the infection with high probability.

Theorem 2. *Fix a failure probability $\delta \in (0,1)$ and an arrival time $k \in \mathbb{N}$. Suppose that each node v has contact probability $q_v \leq 1$. There is a $p(\delta,k) \in (0,1]$ such that, if each node v has infection probability $p_v < p(\delta,k)$, then any non-trival policy contains the infection with probability at least $1-\delta$.*

Proof Sketch. Recall that only nodes with infected parents ever enter the frontier. Additionally, recall that at each step either the frontier is empty (and thus the infection is contained) or the tracer stabilizes one node from the frontier. Therefore, by time $t > k$, either the infection is contained or the tracer has stabilized at least $t-k$ nodes with infected parents.

Let $T'_0, T'_1, \dots$ be the transcript of a tree with the given contagion parameters. The proof idea is to show that there is a time $m > k$ at which point with probability at least $1-\delta$, T'_m has fewer than $m-k$ nodes with infected parents. Since T_m has at most as many nodes with infected parents as T'_m, and since the tracer will have stabilized at least $m-k$ nodes with infected parents by time $t = m$, it follows that with probability at least $1-\delta$ the infection is contained by step m. **??** provides the full proof of theorem 2. □

With the above two theorems, we have established that if either the infection probability or the contact probability is below a certain threshold, any policy contains the infection with high probability.

2.2 Conditions Under Which Containment is Unlikely

Here we show that if both the infection probability and contact probability are above a certain threshold for all nodes, then the infection is unlikely to ever be contained.

Theorem 3. *Fix a policy P. Fix $\delta \in (0,1)$ and $k \geq 3$. There exist $p, q < 1$ such that, if for all nodes v $p_v \geq p$ and $q_v \geq q$, with probability at least $1-\delta$, P does not contain the infection.*

Proof Sketch. The proof idea is to first show that, if at any point in time the number of active infections reaches a certain constant B, then with high probability the number of active infections continues to grow exponentially regardless of the tracer's actions. Second, we show that for p and q large enough, with high probability the first B nodes generated through the infection process are infected nodes. We take a union bound over these two events to prove the theorem. **??** provides the full proof of theorem 3. □

Thus, we've shown that there are settings in which any non-trivial policy is likely to contain the infection, as well as settings in which *no* non-trivial policy is likely to contain the infection. However, does it matter which non-trivial policy we employ? The remainder of the paper studies the following question in a variety of different settings.

Question 2. *Is there an instance and a pair of policies $\mathcal{A}_1$ and $\mathcal{A}_2$ so that the probability of containment under $\mathcal{A}_1$ is greater than the probability of containment under $\mathcal{A}_2$?*

The following section proves that yes, there are settings in which different policies result in different probabilities of containment.

3 POLICY CHOICE AFFECTS PROBABILITY OF CONTAINMENT

In this section we show that policy choice affects the probability of containment by providing an instance in which two non-trivial policies have different probabilities of containment. Recall the two policies presented in section 1, ascending-time and descending-time. Ascending-time prioritizes nodes in order of increasing time-of-arrival and descending-time prioritizes nodes in order of decreasing time-of-arrival, as illustrated in **??** . In this section we demonstrate an instance in which descending-time has a strictly higher probability of containment than ascending-time.

In theorem 4 we prove that, for a certain setting of infection parameters, descending-time has a strictly higher probability of containment than ascending-time. Our analysis in this proof is centered around the different choices the two policies make for a particular setting of parameters, illustrated in fig. 3.

Theorem 4. *There is an instance in which two policies have different probabilities of containment.*

Proof. We provide an instance in which ascending-time and descending-time have different probabilities of containment. Let $p = 0.9999985$, let $q = 1$, and let $k = 3$. Consider the instance where for all nodes v, $p_v = p$ and $q_v = q$. For this instance, let $\mathcal{P}_A$ be the probability of containment for ascending-time, and let $\mathcal{P}_D$ be the probability of containment for descending-time. By analyzing the first few nodes each policy queries, we will show that $\mathcal{P}_D > \mathcal{P}_A$.

The infection process is detailed in fig. 3. The intuition for the proof rests on the fact that the left subtree (rooted at b) is initiated one step before the right subtree (rooted at c). If b is infected, containing the infection in the left subtree is nearly hopeless, due to this headstart. However, if c is infected, we may still be able to chase down the infection in the right subtree. An algorithm that queries c first (descending-time) has a shot at containing the infection in c's subtree and, if b is uninfected, successfully contains the infection. An algorithm that queries b first (ascending-time) allows the infection to continue spreading in c, at which point it may be hopeless to contain the infection in the right subtree.

The proof analyzes five different infection status outcomes for the nodes a, b, and c, which partition the space of all outcomes, as summarized in fig. 4. The full analysis is provided in **??**. We define a failure parameter $\delta = .001$. For a node v, E_v is the event that v is infected. For each outcome u, we bound the probability of containment for both policies conditional on outcome u, where $\mathcal{P}_A(u)$ is the probability of containment for ascending-time and $\mathcal{P}_D(u)$ is the probability of containment for descending-time. We then upper bound $\mathcal{P}_A$ and lower bound $\mathcal{P}_D$ by computing the average probability of containment over the outcomes u, weighted by the probability $\Pr(u)$ that the outcome u occurs.

Using the bounds in fig. 4, we can lower bound $\mathcal{P}_D$ and upper bound $\mathcal{P}_A$.

$$\mathcal{P}_D \geq (1-p) + p(1-p)^2 + p^2(1-p)$$
$$\mathcal{P}_A \leq (1-p) + p(1-p)^2 + 2\delta p^2(1-p) + \delta^2 p^3$$

Therefore,

$$\begin{aligned}\mathcal{P}_D - \mathcal{P}_A &\geq p^2(1-p) - 2\delta p^2(1-p) - \delta^2 p^3 \\ &> 4.97 \times 10^{-7}\end{aligned}$$

Thus there is an instance for which descending-time has a strictly higher probability of containment than ascending-time. □

4 TIME-OF-ARRIVAL HEURISTICS

As introduced in the previous section, there are two obvious policies for prioritizing nodes by time-of-arrival. The *descending-time* policy prioritizes nodes in order of descending time-of-arrival, and the *ascending-time* policy prioritizes nodes based on ascending time-of-arrival. The main question is, which policy has a higher probability of containment?

The previous section provided a single setting of contagion parameters p and q in which descending-time provably has a higher probability of containment than ascending-time. In this section we explore the performance of both policies across a wide range of contagion parameters via computational experiments. Our main finding is that each of the two policies commands a substantial region of the parameter space in which it enjoys a higher probability of containment than the other. These computational experiments extend the conclusions of the previous section by demonstrating numerous instances in which policy choice affects the probability of containment. Moreover, since neither policy has the higher probability of containment in all instances, any comparison between these two policies must take into account the contagion parameters.

Finally, our results qualitatively suggest that a trade-off between p and q may define the boundaries of these regions of dominance. Further characterizing these regions is an intriguing direction for future work.

Simulation overview. Our computational experiments focus on the simple setting where every node is governed by the same infection parameters $p, q \in [0, 1]$. As a result, nodes differ only by time-of-arrival. Our simulation implements the model described in section 1, where $k = 3$, D_p is the constant distribution on p, and D_q is the constant distribution on q. Therefore an instance is defined by the pair (p, q). As defined in the model from section 1, the infection is not contained if an infinite number of nodes become infected. Since checking this condition is intractable, for the purposes of our simulation we redefine containment in terms of a constant $Z_C = 10$; the infection is not contained if more than Z_C nodes are active and infected.[2] If the tracer stabilizes every infected node before this threshold is reached, then the infection is contained.

Even with such a constraint, the infection tree may grow quite large before either of the above two terminating conditions is reached. To manage the size of the tree, our implementation only includes nodes in the tree that could at some point enter the frontier, which are exactly the infected nodes and their children. Even so,

[2]The process for choosing Z_C is described in **??**.

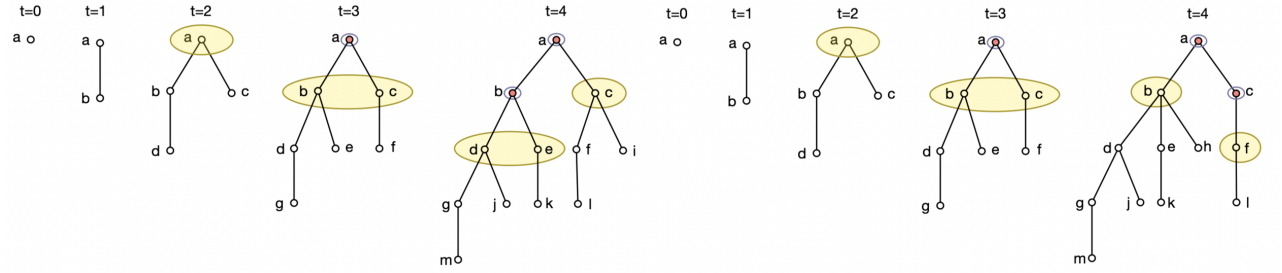

(a) At $t = 4$, ascending-time queries b. **(b) At $t = 4$, descending-time queries c.**

Figure 3: The example shows ascending-time and descending-time operating on the same transcript in a setting where all nodes have the same infection probability $p = .9999985$ and contact probability $q = 1$. The infection process begins at $t = 0$ and runs uninhibited for three steps. At $t = 3$, both policies query a, and since a is infected its children b and c both join the frontier. Node b has time-of-arrival $\tau_b = 1$ while node c has time-of-arrival $\tau_c = 2$. Therefore at $t = 4$ ascending-time queries b while descending-time queries c. We analyze these two choices in theorem 4 to prove that, for this setting of parameters p and q, descending-time has a strictly higher probability of containment.

u	$\Pr(u)$	$\mathcal{P}_A(u)$	$\mathcal{P}_D(u)$
$\neg E_a$	$1-p$	1	1
$E_a \wedge \neg E_b \wedge \neg E_c$	$p(1-p)^2$	1	1
$E_a \wedge \neg E_b \wedge E_c$	$p^2(1-p)$	$\leq \delta$	1
$E_a \wedge E_b \wedge \neg E_c$	$p^2(1-p)$	$\leq \delta$	$\leq \delta$
$E_a \wedge E_b \wedge E_c$	p^3	$\leq \delta^2$	$\leq \delta$

Figure 4: By considering the infection status outcomes for the first few nodes in the tree, we can bound the probabilities of containment for ascending-time and descending-time. Here u is an outcome describing the infection status of nodes a, b, c and $\Pr(u)$ is the probability that outcome u occurs. Given the outcome u, $\mathcal{P}_A(u)$ is the probability that ascending-time contains the infection, and $\mathcal{P}_D(u)$ is the probability that descending-time contains the infection.

there are still instances in which this tree could grow very large.[3] Therefore, we limit the size of the tree to $Z_T = 1000$ nodes. We will argue below that this has only negligible effects on the computational results.

A *trial* is a single run of the simulation and terminates in one of the following three states:

(i) *The infection is contained:* All infected nodes are stabilized.
(ii) *The infection is not contained:* The number of active infections exceeds Z_C.
(iii) *The trial did not converge:* The number of nodes in the tree exceeds Z_T before either (i) or (ii) occur.

Nearly all trials terminate in states (i) or (ii); out of the approximately 2.88×10^{11} trials run in total for all the computational experiments we present, only 87 trials terminated in state (iii). Given the miniscule fraction of trials that terminated in state (iii), it seems unlikely that increasing Z_T would affect our results.

For a fixed instance (p, q), we define a policy's probability of containment to be the probability that a trial terminates in state (i). A policy's *observed* probability of containment for a given instance over a series of trials is the fraction of trials which terminate in state (i). Our goal in the computational experiments that follow is to determine, for a given instance, which policy has the higher probability of containment based on each policy's observed probability of containment over a series of trials.

Computational experiments, part 1. Our computational experiments[4] compare the probabilities of containment for descending-time and ascending-time as the infection probability p and the contact probability q range from .01 to 1 in increments of .01. Since there are 100 possible values each of p and q, this makes for 10^4 instances in total. For each instance (p, q), we run $N = 7.5 \times 10^6$ trials for each policy. The observed probabilities of containment across the entire parameter space are displayed in plots (a) and (b) in fig. 5.

Results, part 1. Plots in (a) and (b) in fig. 5 summarize the results of our experiments. Intuitively, as p and q increase, the observed probability of containment decreases for both policies. For example, when at least one of p or q is at most 0.4, the smallest observed probability of containment for ascending-time is 0.875 and for descending-time is 0.902. The observed probability of containment drops off quickly when both p and q are large. When $p = q = .9$, the observed probability of containment for ascending-time is 0.231 and for descending-time is 0.293; when $p = q = .95$, the observed probability of containment for descending-time is 0.108 and for descending-time is 0.148.

In comparing the two policies, we find that as p and q both grow large, descending-time has the higher observed probability of containment in many instances, and often by large margins. Even in instances where ascending-time exhibits the higher observed probability of containment, the margins are too small to make any claims of statistical significance. It is therefore consistent with this first round of computational experiments that descending-time

[3] See ??.

[4] All computational experiments were run on a server with 144 Intel Xeon Gold 6254 CPUs and 1.5TB RAM running Ubuntu 20.04.4 LTS (Focal Fossa).

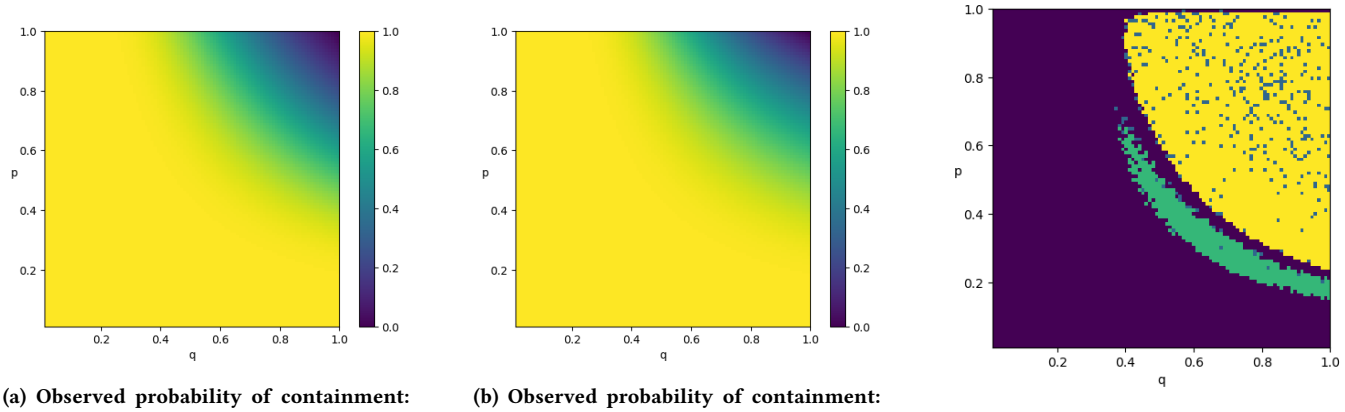

(a) Observed probability of containment: ascending-time

(b) Observed probability of containment: descending-time

(c) Regions of dominance

Figure 5: Plots (a) and (b) show the observed probability of containment for ascending-time and descending-time, respectively, from our first round of computational experiments. Both policies contain the infection with high probability when at least one of p or q is low. When p and q are both large, the observed probability of containment decreases. Plot (c) shows regions where the two policies have significantly different probabilities of containment. In the yellow region, descending-time is significantly better, and in the green region time ascending-time is significantly better. The purple and blue points indicate instances where we do not have sufficient confidence to state which policy is better.

might always be the better policy. Could there be any instance where ascending-time has the higher probability of containment?

Computational experiments, part 2. To investigate this further, we run a second round of computational experiments, in which we focus solely on instances where the absolute difference between the observed probabilities of containment for the two policies is above a fixed threshold. For these instances, we run a second round of trials. Based solely on the trials run in the second round, we compute the observed probabilities of containment for the two policies. We say that a policy *dominates* an instance if we determine with confidence at least $1/2$ that it has a higher probability of containment than the other. Instances where ascending-time dominates are colored green, and instances where descending-time dominates are colored yellow. All other instances are colored purple. Details of the second round of computational experiments are described in **??**.

Finally, to explore whether it is possible to achieve higher confidence bounds, we run a series of 1.5×10^9 trials for each policy for instance parameters $p = .19$ and $q = 1$, since this emerged from the first set of computational experiments as a natural candidate instance where ascending-time might be the better policy.

Results, part 2. Plot (c) in fig. 5 summarizes the results of the second round of trials. As shown in the plot, descending-time dominates a substantial region of the parameter space where p and q are both large, while ascending-time dominates a band directly below. There are 534 instances where ascending-time dominates and 3129 instances where descending-time dominates. We make no claims for the remaining 6538 instances. As a note, one can qualitatively observe these regions of dominance after running far fewer trials, however a large number of trials is necessary for our confidence guarantees.

Given the confidence bounds for plot (c), we can make some further conclusions. First, we know that in any 5×5 square of green instances, with confidence at least $1 - 2^{-25}$, there is at least one instance where $\mathcal{P}_A(p, q) > \mathcal{P}_D(p, q)$. Likewise, and with the same confidence, in any 5×5 square of yellow instances there is at least one instance where $\mathcal{P}_D(p, q) > \mathcal{P}_A(p, q)$. Therefore, with high probability there is a yellow cresent of instances where descending-time has the higher probability of containment, as well as a green crescent of instances directly below where ascending-time has the higher probability of containment. We discuss some compelling directions for future work related to characterizing this border region in section 6.

Finally, while the guarantees for a individual instance in plot (c) only hold with confidence $1/2$, higher confidence guarantees in these regions are also possible. From our series of 1.5×10^9 trials for instance parameters $p = .19$ and $q = 1$ we found that $\mathcal{P}_A(.19, 1) > \mathcal{P}_D(.19, 1)$ with confidence at least $1 - 10^{-10}$.

Comparison to other time-based policies. Ascending-time and descending-time are two natural policies one might consider in practice, and we would like to understand how they compare to other time-based heuristics. Since there are countless ways to prioritize nodes, how can we search the space to see whether these policies are near-optimal?

We implement a heuristic search over policies that order nodes by time-of-arrival, which we detail in **??**. Our search finds no other policy which far outperforms both ascending-time and descending-time. To be clear, this null result does not preclude the existence of such policies, which perhaps could be identified via more advanced techniques. Rather, the fact that these two policies perform on par with policies found via a broad, heuristic search suggests that ascending-time and descending-time are reasonable policies to analyze and study.

5 BEYOND TIME-OF-ARRIVAL HEURISTICS

In all the instances we have seen so far, every node is governed by the same contagion parameters p and q. In this section we explore instances where different nodes may have different contagion

parameters. This opens the door to many more potential policies. Whereas nodes in previous sections were only distinguishable to the tracer based on their time-of-arrival, now a node's contagion parameters may play a factor in its prioritization. The main question is, how should we prioritize nodes, given these three dimensions p_v, q_v, and t_v?

We explore three simple policies for prioritizing nodes: by descreasing infection probability p_v, by decreasing contact probability q_v, and by decreasing time-of-arrival t_v, which is exactly the descending-time policy. We compare the performance of these three policies across a range of instances via computational experiments. We say that a policy *dominates* an instance if, of the three policies, it has the highest probability of containment. We find that each policy dominates a large region of the instance space, which further supports our conclusions from section 4 that the performance of different policies varies across instances.

Of course, each of the policies we consider prioritizes nodes based on only a single dimension, and there are clearly numerous other policies to consider. A compelling direction for future work is to analyze policies which prioritize nodes based on multiple dimensions.

Computational experiments. We compare the performance of the three policies across a range of instances. Recall that we are studying a mathematical process with a genuine probability of containment which is different from the average observed over a series of trials. The goal of our computational experiments is to determine which policy has the highest probability of containment in each instance, based on each policy's observed performance over a series of trials. A node v has infection probability $p_v \sim D_q$, where D_q is the uniform distribution on $[p_{\min}, 1)$, and contact probability $q_v \sim D_q$, where D_q is the uniform distribution on $[q_{\min}, 1)$. Otherwise, the trials follow the same process as in section 4. To generate our space of instances, we let $p_{\min}$ and $q_{\min}$ range independently from 0 to 1 in increments of .01. This results in a total of 10,201 instances.

For each instance, we run 5×10^5 trials for each policy. Via **??**, we compute the confidence that the policy with the highest observed probability of containment is also the policy with the highest (genuine) probability of containment. If the confidence is at least $1/2$, the instance is marked with the color of the dominating policy.

Results. Figure 6 describes our results. Recall that these computational experiments involve infection and contact parameters drawn from uniform distributions on $[p_{\min}, 1)$ and $[q_{\min}, 1)$, respectively, while in section 4 all nodes in a given instance are governed by the same parameters p and q. First, we see that when the infection and contact probabilities for all nodes are large, which happens when $p_{\min}$ and $q_{\min}$ are both large, descending-time dominates.

We do not yet have a mathematical understanding of these regions or an explanation for when one policy might dominate another, however we do draw one insight from these results. Consider the green region where $p_{\min}$ is small and $q_{\min}$ is large. Since $D_q = \text{Unif}[q_{\min}, 1)$, a large value of $q_{\min}$ implies that D_q has small variance. Meanwhile, since $D_p = \text{Unif}[p_{\min}, 1)$, a small value of $p_{\min}$ implies that D_p has large variance. As we see in the plot, in this region the policy which prioritizes by probability of infection p_v dominates. A similar phenomenon occurs in the yellow region where $q_{\min}$ is small and $p_{\min}$ is large. In this region the policy which prioritizes by contact probability q_v dominates. The intuition that we take away from this plot is that if there are very large differences in the variances of D_p and D_q, it is better to prioritize by the parameter with the larger variance. There are numerous directions for future work in this areas, which we outline in section 6.

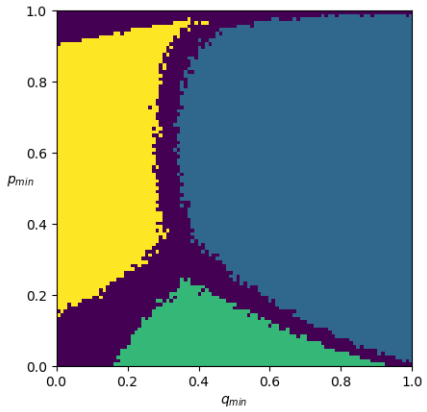

Figure 6: The above figure illustrates results from computational experiments where $p_v \sim \text{Unif}[p_{\min}, 1)$ and $q_v \sim \text{Unif}[q_{\min}, 1)$. Each pixel corresponds to an instance defined by $(p_{\min}, q_{\min})$. We compare the probabilities of containment for three policies: prioritizing by p_v, prioritizing by q_v, and descending-time. We say that a policy *dominates* if its probability of containment is signficantly higher than that of the other two policies. Each pixel is colored to indicate which of the three policies, prioritizing by p_v (green), prioritizing by q_v (yellow), or descending-time (blue) dominates. Purple indicates areas of no confidence.

6 CONCLUSION

In this paper we show, via both theoretical bounds and computational experiments, that policy choice affects probability of containment in many different settings and that the best policy to deploy may depend on the parameters of the contagion.

One interesting avenue for future work would be to prove that there is an instance where ascending-time has a higher probability of containment than descending-time, complementing the result for descending-time in section 3. Characterizing the curve in the $p - q$ parameter space of fig. 5 plot (c) is another interesting direction, as is exploring other policies for prioritizing nodes beyond time-based methods.

ACKNOWLEDGMENTS

Michela Meister was supported by the Department of Defense (DoD) through the National Defense Science & Engineering Graduate (NDSEG) Fellowship Program. Jon Kleinberg was supported in part by a Simons Investigator Award, a Vannevar Bush Faculty Fellowship, MURI grant W911NF-19-0217, AFOSR grant FA9550-19-1-0183, ARO grant W911NF19-1-0057, and a grant from the MacArthur Foundation.

ETHICS STATEMENT

This paper studies a model of contact tracing and aims to understand when different contact tracing policies are likely to contain an outbreak of an infection. The results of this paper are comprised of mathematical proofs and computational simulations. No datasets were used in the work for this paper.

This paper takes a theoretical approach, and we do not consider all the factors that would need to be addressed if one were to implement such policies in the real world. Were such policies implemented in the real world, we think it would be important to ensure that whatever prioritization policy used is implemented equitably. For example, a policy which prioritizes individuals by recency of infection may not explore the contacts of individiuals infected early on in the infection process. A team implementing this policy would need to ensure that such prioritization would not result in inequitable access for the given population. This would be particularly important if traced nodes receive other benefits, such as medical treatment.

REFERENCES

[1] Benjamin Armbruster and Margaret L Brandeau. 2007. Contact tracing to control infectious disease: when enough is enough. *Health care management science* 10, 4 (2007), 341–355.

[2] Benjamin Armbruster and Margaret L Brandeau. 2007. Who do you know? A simulation study of infectious disease control through contact tracing. In *Proceedings of the 2007 Western Multiconference on Computer Simulation*. 79–85.

[3] National Tuberculosis Controllers Association, Centers for Disease Control, Prevention (CDC, and others). 2005. Guidelines for the investigation of contacts of persons with infectious tuberculosis. Recommendations from the National Tuberculosis Controllers Association and CDC. *MMWR. Recommendations and reports: Morbidity and mortality weekly report. Recommendations and reports* 54, RR-15 (2005), 1–47.

[4] Omri Ben-Eliezer, Elchanan Mossel, and Madhu Sudan. 2021. Information Spread with Error Correction. *arXiv preprint arXiv:2107.06362* (2021).

[5] KTD Eames. 2007. Contact tracing strategies in heterogeneous populations. *Epidemiology & Infection* 135, 3 (2007), 443–454.

[6] Ken TD Eames and Matt J Keeling. 2003. Contact tracing and disease control. *Proceedings of the Royal Society of London. Series B: Biological Sciences* 270, 1533 (2003), 2565–2571.

[7] Christophe Fraser, Steven Riley, Roy M Anderson, and Neil M Ferguson. 2004. Factors that make an infectious disease outbreak controllable. *Proceedings of the National Academy of Sciences* 101, 16 (2004), 6146–6151.

[8] Joel Hellewell, Sam Abbott, Amy Gimma, Nikos I Bosse, Christopher I Jarvis, Timothy W Russell, James D Munday, Adam J Kucharski, W John Edmunds, Fiona Sun, et al. 2020. Feasibility of controlling COVID-19 outbreaks by isolation of cases and contacts. *The Lancet Global Health* 8, 4 (2020), e488–e496.

[9] Herbert W Hethcote, James A Yorke, and Annett Nold. 1982. Gonorrhea modeling: a comparison of control methods. *Mathematical Biosciences* 58, 1 (1982), 93–109.

[10] Leslie Pack Kaelbling, Michael L Littman, and Anthony R Cassandra. 1998. Planning and acting in partially observable stochastic domains. *Artificial intelligence* 101, 1-2 (1998), 99–134.

[11] Satwik Kansal and Brendan Martin. [n.d.]. Reinforcement Q-Learning from Scratch in Python with OpenAI Gym. https://www.learndatasci.com/tutorials/reinforcement-q-learning-scratch-python-openai-gym/.

[12] Edward H Kaplan, David L Craft, and Lawrence M Wein. 2002. Emergency response to a smallpox attack: the case for mass vaccination. *Proceedings of the National Academy of Sciences* 99, 16 (2002), 10935–10940.

[13] Edward H Kaplan, David L Craft, and Lawrence M Wein. 2003. Analyzing bioterror response logistics: the case of smallpox. *Mathematical biosciences* 185, 1 (2003), 33–72.

[14] Don Klinkenberg, Christophe Fraser, and Hans Heesterbeek. 2006. The effectiveness of contact tracing in emerging epidemics. *PloS one* 1, 1 (2006), e12.

[15] Mirjam E Kretzschmar, Ganna Rozhnova, Martin CJ Bootsma, Michiel van Boven, Janneke HHM van de Wijgert, and Marc JM Bonten. 2020. Impact of delays on effectiveness of contact tracing strategies for COVID-19: a modelling study. *The Lancet Public Health* 5, 8 (2020), e452–e459.

[16] Kin On Kwok, Arthur Tang, Vivian WI Wei, Woo Hyun Park, Eng Kiong Yeoh, and Steven Riley. 2019. Epidemic models of contact tracing: systematic review of transmission studies of severe acute respiratory syndrome and Middle East respiratory syndrome. *Computational and structural biotechnology journal* 17 (2019), 186–194.

[17] George Z Li, Arash Haddadan, Ann Li, Madhav Marathe, Aravind Srinivasan, Anil Vullikanti, and Zeyu Zhao. 2022. Theoretical Models and Preliminary Results for Contact Tracing and Isolation. In *Proceedings of the 21st International Conference on Autonomous Agents and Multiagent Systems*.

[18] Michael L Littman. 2009. A tutorial on partially observable Markov decision processes. *Journal of Mathematical Psychology* 53, 3 (2009), 119–125.

[19] Maia Martcheva. 2015. *An introduction to mathematical epidemiology*. Vol. 61. Springer.

[20] Michela Meister and Jon Kleinberg. 2021. Optimizing the order of actions in contact tracing. *arXiv preprint arXiv:2107.09803* (2021).

[21] Michael Mitzenmacher and Eli Upfal. 2017. *Probability and computing: Randomization and probabilistic techniques in algorithms and data analysis*. Cambridge university press.

[22] Mehryar Mohri, Afshin Rostamizadeh, and Ameet Talwalkar. 2018. *Foundations of machine learning*. MIT press.

[23] Johannes Müller and Mirjam Kretzschmar. 2021. Contact tracing–Old models and new challenges. *Infectious Disease Modelling* 6 (2021), 222–231.

[24] Johannes Müller, Mirjam Kretzschmar, and Klaus Dietz. 2000. Contact tracing in stochastic and deterministic epidemic models. *Mathematical biosciences* 164, 1 (2000), 39–64.

[25] STD National Center for HIV/AIDS, Viral Hepatitis and TB Prevention (CDC). 2008. Recommendations for partner services programs for HIV infection, syphilis, gonorrhea, and chlamydial infection. *MMWR. Recommendations and reports: Morbidity and mortality weekly report. Recommendations and reports* 57, RR-9 (2008), 1–92.

[26] Ronald Parr and Stuart Russell. 1995. Approximating optimal policies for partially observable stochastic domains. In *IJCAI*, Vol. 95. Citeseer, 1088–1094.

[27] Martin L Puterman. 2014. *Markov decision processes: discrete stochastic dynamic programming*. John Wiley & Sons.

[28] Guni Sharon, James Ault, Peter Stone, Varun Kompella, and Roberto Capobianco. 2021. Multiagent Epidemiologic Inference through Realtime Contact Tracing. In *Proceedings of the 20th International Conference on Autonomous Agents and MultiAgent Systems* (Virtual Event, United Kingdom) *(AAMAS '21)*. International Foundation for Autonomous Agents and Multiagent Systems, Richland, SC, 1182–1190.

[29] Kimberly D Spencer, Christina L Chung, Alison Stargel, Alvin Shultz, Phoebe G Thorpe, Marion W Carter, Melanie M Taylor, Mary McFarlane, Dale Rose, Margaret A Honein, et al. 2021. COVID-19 case investigation and contact tracing efforts from health departments—United States, June 25–July 24, 2020. *Morbidity and Mortality Weekly Report* 70, 3 (2021), 83.

[30] Yuan Tian, Nathaniel D Osgood, Assaad Al-Azem, and Vernon H Hoeppner. 2013. Evaluating the effectiveness of contact tracing on tuberculosis outcomes in Saskatchewan using individual-based modeling. *Health education & behavior* 40, 1_suppl (2013), 98S–110S.

[31] Andres I Vecino-Ortiz, Juliana Villanueva Congote, Silvana Zapata Bedoya, and Zulma M Cucunuba. 2021. Impact of contact tracing on COVID-19 mortality: An impact evaluation using surveillance data from Colombia. *Plos one* 16, 3 (2021), e0246987.

[32] Christopher JCH Watkins and Peter Dayan. 1992. Q-learning. *Machine learning* 8, 3 (1992), 279–292.