

# Risk-Averse Antibiotics Time Machine Problem

Deniz Tuncer

Industrial Engineering Program, Sabancı University, 34956 Istanbul, Turkey, dtuncer@sabanciuniv.edu

Burak Kocuk

Industrial Engineering Program, Sabancı University, 34956 Istanbul, Turkey, burak.kocuk@sabanciuniv.edu

Antibiotic resistance, which is a serious healthcare issue, emerges due to uncontrolled and repeated antibiotic use that causes bacteria to mutate and develop resistance to antibiotics. The Antibiotics Time Machine Problem aims to come up with treatment plans that maximize the probability of reversing these mutations. Motivated by the severity of the problem, we develop a risk-averse approach and formulate a scenario-based mixed-integer linear program with a conditional value-at-risk objective function. We propose a risk-averse scenario batch decomposition algorithm that partitions the scenarios into manageable risk-averse subproblems, enabling the construction of lower and upper bounds. We develop several algorithmic enhancements in the form of stronger no-good cuts and symmetry breaking constraints in addition to scenario regrouping and warm starting. We conduct extensive computational experiments for static and dynamic versions of the problem on a real dataset and demonstrate the effectiveness of our approach. Our results suggest that risk-averse solutions can achieve significantly better worst-case performance compared to risk-neutral solutions with a slight decrease in terms of the average performance, especially for the dynamic version. Although our methodology is presented in the context of the Antibiotics Time Machine Problem, it can be adapted to other risk-averse problem settings in which the decision variables come from special ordered sets of type one.

*Key words:* Stochastic programming, mixed-integer programming, conditional value-at-risk, antibiotic resistance.

---

## 1. Introduction and Literature Review

Antibiotic resistance is a serious concern in modern medicine. Misuse or successive administration of antibiotics may cause mutations of the bacteria, which facilitate the resistance to the known antibiotics. Due to their immunization to the antibiotics, it may not be possible to cure the diseases caused by the antibiotic-resistant bacteria. 2.8 million antibiotic-resistant infections occurred in the US, and more than 35 thousand people lost their lives in 2019 (Murray et al. 2022) due to antibiotic resistance. A study on Chinese healthcare system (Zhang et al. 2024) shows that about 600 thousand deaths have occurred due to bacterial antimicrobial resistance. The same study also shows that in China, 173 million unreasonable prescriptions of antibiotics are given each year, which also contributes to the antibiotic resistance. Adverse effects of some overprescribed antibiotics, such as a widely used antibiotic, Amoxicillin-clavulanic acid is shown to cause life threatening adverse

effects (Chang and Schiano 2007), which further suggests that antibiotic administration should be done with utmost care.

The antibiotics that humankind have developed during the history are in danger of becoming obsolete. Therefore, we need new strategies to combat antibiotic resistance. While developing new antibiotics seem to be the obvious solution, research and development process of a new antibiotic is quite costly. Even if the necessary funding is found, the process itself requires extensive human resource and the time it takes to develop and start using the drug is unpredictable (Ventola 2015). A recent study by Liu et al. (2024) shows that only 18 novel antibiotics have been approved since 2014. Since development of new antibiotics is unlikely to happen soon, as an interim solution, researchers have proposed the method of repurposing drugs. Drug repurposing aims to break the antibiotic resistance by co-administering the antibiotics with non-antibiotic drugs (Brown 2015). This repurposing idea may increase the lifespan of the existing antibiotics while the new antibiotics are being developed. While drug repurposing seems to be the recipe for slowing down the antibiotic resistance, there exists concerns about the future efficiency of the approach, since it may cause further antibiotic resistance that we may not be able to cure (Talat et al. 2022).

In search of a remedy for the antibiotic resistance, some of the recent works on the problem focus on artificial intelligence (AI) and deep learning-based methods (Liu et al. 2024). The study by Liu et al. (2023) aims to create new antibiotics by utilizing deep learning methods to develop molecules that have activity against the *Acinetobacter baumannii* bacterium. (Wong et al. 2023) construct explainable graphs using deep learning methods to predict the compounds that might have high antibiotic activity. Apart from drug development, AI and deep learning methods are utilized to predict the antibiotic resistance of the bacteria that cause the infections (Zagajewski et al. 2023, Ding et al. 2024). These methods are also utilized for identifying the drugs that may be eligible for drug repurposing (Aggarwal et al. 2024), which can be a future research direction for use of AI for antibiotic resistance problem.

It is found out that when some antibiotics are applied sequentially, their cross effect may change the probability of the bacteria developing antibiotic resistance, called *collateral sensitivity* (Imamovic and Sommer 2013, Nichol et al. 2015). Although there are opposing ideas that drug sequencing is unlikely to prevent antibiotic resistance (Bergstrom et al. 2004, van Duijn et al. 2018), some studies claim that the collateral sensitivity can eliminate antibiotic resistance (Tyers and Wright 2019, Maltas and Wood 2019, 2021). Another study (Aulin et al. 2021) shows that the order of drug administration is crucial to the efficiency of the collateral sensitivity based treatments.

A further complicating factor in these studies is the issue of quantifying the antibiotic resistance phenomenon (Mira et al. 2015). Typically, the effect of an antibiotic can be modeled as a Markov chain, where the states correspond to the genotypes of the bacterium identified by a string of alleles

and probabilities are computed as a function of the growth rates measurements of these genotypes under the administration of the antibiotic. In addition to the stochasticity of an antibiotic administration, another source of randomness arises due to the fact that the growth rates measurements, and consequently, the transition probability matrices governing the Markov chain can be uncertain as well (Mira et al. 2017).

The Antibiotics Time Machine Problem is an NP-Hard combinatorial optimization problem (Tran and Yang 2017) that aims to find the best treatment plan of finite length using a set of given antibiotics so that the probability of reversing the antibiotic resistance is maximized. The first paper to tackle this problem is Mira et al. (2015), in which the authors use a complete enumeration technique to find an optimal sequence. Recently, a compact-size mixed-integer linear programming (MILP) formulation is proposed in Kocuk (2022) that works much more efficiently. We will refer to the version studied in these three papers as *deterministic* since the probability matrices are computed using a single growth rate measurement and *static* since the antibiotic sequence is determined at the beginning of the treatment. Later, the work of Kocuk (2022) is extended to the cases of *stochastic* and *dynamic* versions of the Antibiotics Time Machine Problem in Mesüm et al. (2024). Here, the stochastic version refers to the case where multiple growth rate measurements are available and the probability matrices are modeled as random variables whereas the dynamic version refers to the decision policies in which the treatment plan is adaptive to the observed state transitions. It is proven in Mesüm et al. (2024) that the stochastic static version of the problem can be solved as a special instance of the deterministic static version, which can be solved as an MILP. On the other hand, the stochastic dynamic version of the problem can be solved as a special instance of the deterministic dynamic version, which can be solved using a dynamic program (DP).

Although the work by Mesüm et al. (2024) presents the first attempt to solve the Antibiotics Time Machine Problem under uncertainty, it assumes that the objective function is the maximization of the *expected* probability of reversing the antibiotic resistance. However, this *risk-neutral* approach may not be suitable for a healthcare application due to its sensitive nature. Therefore, in this paper, we study a *risk-averse* approach, in which we maximize the expected probability of the worst  $\alpha$  fraction of the realizations akin to a Conditional-Value-at-Risk (CVaR)-type objective function. As opposed to the expectation, which is a linear function, CVaR is not a linear function, therefore, the theoretical results from Mesüm et al. (2024) are not applicable. It turns out that both static and dynamic versions of the risk-averse Antibiotics Time Machine Problem are quite challenging computationally.

In this paper, we study scenario-based MILP formulations for both static and dynamic versions of the risk-averse Antibiotics Time Machine Problem for the first time. Since these MILPs do not

scale well with the number of scenarios, we develop a risk-averse scenario batch decomposition algorithm, originally developed for risk-neutral (Ahmed 2013) and risk-averse problems (Deng et al. 2018). As opposed to these papers, which utilize single-scenario subproblems, we use batch scenario subproblems in order to obtain stronger dual bounds. To further improve the computational efficiency of the standard algorithm, we derive two types of no-good cuts that exploit problem structure and solution symmetry. In addition, we introduce scenario regrouping and warm start techniques that help accelerate the decomposition algorithm. Although these contributions to the literature are presented in the context of the Antibiotics Time Machine Problem, most of them are applicable for risk-averse optimization problems in which the decision variables come from special ordered sets of type one.

Our extensive computational experiments conducted using real growth rate measurements (Mira et al. 2017) give interesting insights. As expected, switching from the risk-neutral solution to the risk-averse solution brings considerable gains from the worst case realizations albeit a loss from the average performance. Interestingly, as the treatment length increases, the amount of gain from the worst case realizations starts to dominate the amount of loss from the average performance. In addition, the benefit of the risk-averse approach is even more pronounced for the dynamic version for which solutions with almost negligible sacrifice from the average performance may increase the worst case realizations considerably.

The rest of the paper is organized as follows: In Section 2, we introduce the deterministic Antibiotics Time Machine Problem. In Sections 3 and 4, we study static and dynamic versions of the risk-averse Antibiotics Time Machine Problem, respectively. We propose decomposition algorithms for their MILP formulations and provide several algorithmic enhancements. In Section 5, we present the results of our extensive computational experiments. Finally, we conclude our paper in Section 6.

## 2. Deterministic Problem


In this section, we present the deterministic version of the Antibiotics Time Machine Problem studied in Mira et al. (2015), Kocuk (2022) as this will make it easier to explain the stochastic versions later on. Suppose that we are interested in a bacterium with  $a$  many alleles. We denote each genotype by the vector  $\mathbf{g} \in \{0, 1\}^a$ , where  $g_l = 0$  (resp.  $g_l = 1$ ) indicates an unmutated (resp. mutated) allele  $l$ . We call the special genotype associated with the vector  $\mathbf{g} = \mathbf{0}$  as the “wild type”, which is the desired state where the bacterium is sensitive to an antibiotic treatment.

Suppose we have  $K$  many antibiotics at hand, each of which comes with a probability transition matrix  $\mathbf{M}^k \in \mathbb{R}^{d \times d}$ ,  $k = 1, \dots, K$ , where  $d = 2^a$  (we explain how these matrices are obtained from experimental data in Appendix A). In particular, the probability of transitioning from genotype  $i$  to genotype  $j$  when antibiotic  $k$  is applied is denoted by  $M_{ij}^k$ . Due to the generally accepted principle

of Strong Selection Weak Mutation (see, e.g., Gillespie (1983), Gillespie (1984)), we will assume that at most one allele can change in each transition. In other words, the probability of a transition between genotype  $\mathbf{g}$  and genotype  $\mathbf{g}'$  is zero if  $\|\mathbf{g} - \mathbf{g}'\|_1 > 1$ .

Let us illustrate the setup with an example in Figure 1 in which there are  $d = 2^3$  genotypes and  $K = 2$  antibiotics represented by different types of arcs. An arc indicates that there is a positive probability of going from a state to another under that antibiotic. For example, starting from genotype 010, if we apply antibiotic 1 (solid line), the stochastic process can go to the genotype 000 or genotype 011. If the current state is 101 and we apply antibiotic 2 (dashed line), the next state would be 111.

**Figure 1** Illustration of an Antibiotics Time Machine Problem instance with  $a = 3$  alleles and  $K = 2$  antibiotics.



The goal in the Antibiotics Time Machine Problem is to come up with a treatment plan that maximizes the probability of reaching the wild type starting from a given initial genotype with a predetermined number ( $N$ ) of antibiotic applications. We will consider two versions of the problem depending on the type of plans allowed: static and dynamic. The solution of the static version is a sequence of antibiotics determined at the beginning of the planning horizon. For the instance in Figure 1, an example sequence is 1-2-2, which gives the probability of reaching the wild type  $M_{111,011}^1 \cdot M_{011,001}^2 \cdot M_{001,000}^2$  if the initial genotype is 111. In the dynamic version, the solution is expressed by a policy, where we decide which antibiotic to apply based on the genotype at each decision epoch. Note that the dynamic version requires intermediate observations of state transitions. For the instance in Figure 1, an example policy would be to apply antibiotic  $k = 1 + g_1$  if the current state is  $\mathbf{g}$ , which gives the probability of reaching the wild type  $M_{111,011}^1 \cdot M_{011,001}^2 \cdot M_{001,000}^2 + M_{111,110}^1 \cdot M_{110,100}^1 \cdot M_{100,000}^1$  if the initial genotype is 111.

In the rest of the paper, we will consider the stochastic version of the Antibiotics Time Machine Problem in which the matrices  $\mathbf{M}^k$  are uncertain. In particular, we will assume that  $\mathbf{M}^k$  is a discrete random variable that is equally likely to take value from the set of all possible antibiotic matrices  $\{\mathbf{M}^{k,h} \in \mathbb{R}^{d \times d} : h \in \mathcal{S}\}$ , where  $\mathcal{S}$  is an index set (see Appendix A for details). In reality,

the cardinality of  $\mathcal{S}$  can be quite large (e.g.,  $12^{16}$ , see, Mesüm et al. (2024)). Instead, we will employ Sample Average Approximation (SAA) and take a random sample  $\mathcal{H} \subseteq \mathcal{S}^K$  such that  $|\mathcal{H}| \ll |\mathcal{S}|^K$ .

### 3. Static Stochastic Problem

In this section, we will assume that we are given a stochastic Antibiotics Time Machine Problem instance with  $d = 2^a$  many alleles,  $K$  many antibiotics and a scenario sample  $\mathcal{H}$ . We present how to formulate and solve the risk-neutral and risk-averse versions of this problem in Sections 3.1 and 3.2, respectively. Then, we propose a scenario decomposition algorithm in Section 3.3 and several algorithmic enhancements to increase its computational efficiency in Section 3.4.

#### 3.1. Risk-Neutral Version

In this version of the problem, we assume that the objective function maximizes the *expected probability* of reaching the wild type in  $N$  steps given an initial genotype. In order to formulate this problem, we first introduce some notation: Let  $\mathbf{r} \in \mathbb{R}^d$  and  $\mathbf{q} \in \mathbb{R}^d$  denote the standard unit vectors which correspond to the initial distribution and the desired distribution, respectively (we will use a genotype, a binary vector in  $\{0, 1\}^a$ , and its ID, a positive integer in  $\{1, \dots, 2^a\}$ , interchangeably throughout the paper). We define a binary decision variable  $x_{n,k}$  that takes the value 1 if antibiotic  $k$  is applied at step  $n$ , and 0 otherwise. Then, the static risk-neutral optimization model can be formulated as

$$\max_{\mathbf{x} \in \mathcal{X}} \frac{1}{|\mathcal{H}|} \sum_{h \in \mathcal{H}} \mathbf{r}^\top \underbrace{\left[ \prod_{n=1}^N \left( \sum_{k=1}^K M^{k,h} x_{n,k} \right) \right]}_{:=f_h(\mathbf{x})} \mathbf{q}, \quad (1)$$

where the set  $\mathcal{X} := \{\mathbf{x} \in \{0, 1\}^{N \times K} : \sum_{k=1}^K x_{n,k} = 1, n = 1, \dots, N\}$  ensures that exactly one antibiotic  $k$  is chosen for each position  $n$ . The objective function is calculated for each scenario  $h$  depending on the antibiotic selection by multiplying the associated probability transition matrices, and then, averaged over  $h \in \mathcal{H}$ . Notice that the formulation (1) is a nonlinear integer programming problem, which seems difficult to solve directly. Instead, we will introduce new variables, using which the problem is reformulated as an MILP. For this purpose, let  $\mathbf{u}_n^h$  denote the probability distribution after the  $n$ -th antibiotic is applied under scenario  $h$ . Then, our decision variables satisfy the following recursion:

$$\sum_{k=1}^K (\mathbf{u}_{n-1}^h)^\top M^{k,h} x_{n,k} = (\mathbf{u}_n^h)^\top \quad n = 1, \dots, N. \quad (2)$$

Adapting the deterministic formulation in Kocuk (2022), Mesüm et al. (2024) to this setting, it is possible to reformulate problem (1) exactly as the following MILP:

$$\text{S-RN}(\mathcal{H}): \quad \max_{\mathbf{u}, \mathbf{v}, \mathbf{x} \in \mathcal{X}} \frac{1}{|\mathcal{H}|} \sum_{h \in \mathcal{H}} \mathbf{q}^\top \mathbf{u}_N^h \quad (3a)$$

$$\text{s.t. } \mathbf{u}_0^h = \mathbf{r} \quad h \in \mathcal{H} \quad (3b)$$

$$\sum_{k=1}^K \mathbf{v}_{n-1,k}^h = \mathbf{u}_{n-1}^h \quad n = 1, \dots, N, h \in \mathcal{H} \quad (3c)$$

$$\sum_{k=1}^K (\mathbf{v}_{n-1,k}^h)^\top \mathbf{M}^{k,h} = (\mathbf{u}_n^h)^\top \quad n = 1, \dots, N, h \in \mathcal{H} \quad (3d)$$

$$\mathbf{e}^\top \mathbf{v}_{n-1,k}^h = x_{n,k} \quad n = 1, \dots, N, k = 1, \dots, K, h \in \mathcal{H} \quad (3e)$$

$$\mathbf{u}_n^h \in \mathbb{R}_+^d, \mathbf{v}_{n-1,k}^h \in \mathbb{R}_+^d \quad n = 1, \dots, N, k = 1, \dots, K, h \in \mathcal{H}. \quad (3f)$$

In this formulation, the objective function (3a) maximizes the expected probability of reaching the desired state. Constraint (3b) ensures that the initial state distribution is  $\mathbf{r}$ . Constraints (3c)-(3e) linearize the recursive relation (2) between each drug application. Here, variables  $\mathbf{v}_n^h$  are the copy variables that help linearizing this relation and  $\mathbf{e}$  denotes the vector of ones. Finally, constraint (3f) gives the variable domain restrictions.

### 3.2. Risk-Averse Version


While solving the problem with a risk-neutral approach that solely focuses on the average performance seems plausible at first, it is important for a treatment to be successful even for the worst cases due to the sensitive nature of the Antibiotics Time Machine Problem. Therefore, solving the problem in a risk-averse manner would be a more reasonable approach. Our risk measure is CVaR, where we aim to maximize the average of the worst  $\alpha$  fraction of the realizations.

To motivate the relevance of risk-aversion, we compare the solutions obtained from the risk-neutral and risk-averse approaches with the initial genotype 0001 under  $N = 5$ . In this case, the average performances of risk-neutral and risk-averse solutions are 0.55 and 0.52, respectively, which are relatively close. However, the average of the worst 10% of the realizations is 0.00 and 0.31 for risk-neutral and risk-averse solutions, respectively, resulting in a stark difference. To better compare the two approaches, we also provide the histograms corresponding to the frequency of certain probability ranges of reaching the wild type in Figure 2. We clearly see that the risk-neutral solution has a large number of realizations with low probability of reaching the wild type, which is significantly reduced in the risk-averse solution.

Having motivated the risk-averse approach, let us formally define our risk measure CVaR: Let  $\mathbf{v} \in \mathbb{R}^{|\mathcal{H}|}$  and  $\alpha \in (0, 1]$  such that  $\alpha|\mathcal{H}| \in \mathbb{Z}_+$ . Then, we define  $s_\alpha(\mathbf{v})$  as the average of the smallest  $\alpha|\mathcal{H}|$  entries of  $\mathbf{v}$ . We are now ready to provide the risk-averse formulation of the problem:

$$\max_{\mathbf{x} \in \mathcal{X}} s_\alpha([f_h(\mathbf{x})]_{h \in \mathcal{H}}) \quad (4)$$

Although  $s_\alpha(\cdot)$  is a nonlinear function, it has a polyhedral representation as given in Proposition 1.

**Figure 2** Comparison of risk-neutral and risk-averse solutions for the initial genotype 0001 and  $N = 5$ .

PROPOSITION 1. (Ben-Tal and Nemirovski 2001) Consider the function  $s_\alpha(\mathbf{v})$  defined above. Then,

$$s_\alpha(\mathbf{v}) = \max_{\lambda, \mu} \left\{ \lambda - \frac{1}{\alpha|\mathcal{H}|} \sum_{h \in \mathcal{H}} \mu_h : \lambda - \mu_h \leq v_h, \mu_h \geq 0 \ h \in \mathcal{H} \right\}.$$

Combining formulation (3) and Proposition 1, we provide an MILP formulation for the risk-averse version of the problem as below:

$$\text{S-RA}_\alpha(\mathcal{H}) : \max_{\mathbf{u}, \mathbf{v}, \mathbf{x} \in \mathcal{X}, \lambda, \mu} \left\{ \lambda - \frac{1}{\alpha|\mathcal{H}|} \sum_{h \in \mathcal{H}} \mu_h : \lambda - \mu_h \leq \mathbf{q}^\top \mathbf{u}_N^h \ h \in \mathcal{H}, \mu_h \geq 0 \ h \in \mathcal{H}, (3b) - (3f) \right\}.$$

Notice that problem  $\text{S-RA}_1(\mathcal{H})$  is equivalent to  $\text{S-RN}(\mathcal{H})$ , therefore, the risk-neutral formulation is a special case of the risk-averse formulation. Therefore, in the remainder of the paper, we will mainly analyze the risk-averse formulation.

### 3.3. Scenario Decomposition Algorithm

It is challenging to solve the formulations introduced above directly using an off-the-shelf MILP solver even for moderate values of  $|\mathcal{H}|$ . Therefore, we resort to the scenario decomposition algorithm, originally proposed for risk-neutral problems in Ahmed (2013). Our decomposition approach equipartitions the scenario set  $\mathcal{H}$  into subsets  $\{\mathcal{H}^p : p = 1, \dots, P\}$  such that  $\alpha|\mathcal{H}^p| \in \mathbb{Z}_+$ , and solves *scenario batch subproblems* as  $\text{S-RA}_\alpha(\mathcal{H}^p)$ . The critical component of the scenario decomposition algorithm is to derive a lower bound (LB) and an upper bound (UB) from the scenario batch subproblems. The result below provides these bounds for the risk-averse setting.

PROPOSITION 2. Let  $\alpha \in (0, 1]$ . Denote the optimal value of problem  $\text{S-RA}_\alpha(\mathcal{H})$  as  $\hat{z}_\mathcal{H}$  and  $\mathbf{x}^{(p)} = \arg \max_{\mathbf{x} \in \mathcal{X}} s_\alpha([f_h(\mathbf{x})]_{h \in \mathcal{H}^p})$ . Then,

$$\underbrace{\max_{p=1, \dots, P} \{s_\alpha([f_h(\mathbf{x}^{(p)})]_{h \in \mathcal{H}})\}}_{LB} \leq \hat{z}_\mathcal{H} \leq \underbrace{\frac{1}{P} \sum_{p=1}^P s_\alpha([f_h(\mathbf{x}^{(p)})]_{h \in \mathcal{H}^p})}_{UB}.$$



PROOF. We first prove the lower bound as follows: Let  $p = 1, \dots, P$ . Since the solution  $\mathbf{x}^{(p)}$  obtained from a batch subproblem S-RA $_{\alpha}(\mathcal{H}^p)$  can be extended with appropriate  $\mathbf{u}$  and  $\mathbf{v}$  values such that it is feasible for S-RA $_{\alpha}(\mathcal{H})$ , we deduce that  $s_{\alpha}([f_h(\mathbf{x}^{(p)})]_{h \in \mathcal{H}})$  is a lower bound on  $\hat{z}_{\mathcal{H}}$ . Consequently, the best lower bound is obtained by taking the maximum over  $p = 1, \dots, P$ .

We next prove the upper bound as follows: Let  $\mathbf{x}^*$  be an optimal solution for S-RA $_{\alpha}(\mathcal{H})$ . Note that  $\mathbf{x}^*$  is also feasible for batch subproblem S-RA $_{\alpha}(\mathcal{H}^p)$ ,  $p = 1, \dots, P$  when extended with appropriate  $\mathbf{u}$  and  $\mathbf{v}$  values. Since  $\mathbf{x}^{(p)}$  is an optimal solution for S-RA $_{\alpha}(\mathcal{H}^p)$ , the relation  $s_{\alpha}([f_h(\mathbf{x}^*)]_{h \in \mathcal{H}^p}) \leq s_{\alpha}([f_h(\mathbf{x}^{(p)})]_{h \in \mathcal{H}^p})$  holds true. The result follows since  $\hat{z}_{\mathcal{H}} = s_{\alpha}([f_h(\mathbf{x}^*)]_{h \in \mathcal{H}}) \leq \frac{1}{P} \sum_{p=1}^P s_{\alpha}([f_h(\mathbf{x}^*)]_{h \in \mathcal{H}^p})$ .  $\square$

We are now able to present our static scenario decomposition scheme in Algorithm 1. We note that since our approach is based on SAA, upon termination with a solution  $\mathbf{x}^*$  obtained from the “in-sample optimization”, we carry out an “out-of-sample” evaluation using a fresh sample scenario set  $\mathcal{H}'$ .

---

**Algorithm 1** Static Scenario Decomposition Algorithm.

---

Set  $t = 1$ ,  $\mathbf{x}^* = 0$ ,  $LB = 0$  and  $UB = 1$ .

**while**  $UB - LB > \epsilon$  and  $t \leq \tau$  **do**

Solve S-RA $_{\alpha}(\mathcal{H}^p)$  in parallel for  $p = 1, \dots, P$ .

If needed, update  $LB$  and  $UB$  as in Proposition 2 and the best feasible solution  $\mathbf{x}^*$ .

Add the following no-good cut to the set  $\mathcal{X}$  for each distinct solution  $\mathbf{x}^{(p)}$  obtained:

$$\sum_{(n,k): x_{n,k}^{(p)}=1} x_{n,k} \leq |N| - 1. \quad (5)$$

Increment  $t = t + 1$ .

Carry out an “out-of-sample” analysis with a new sample  $\mathcal{H}'$  by computing  $s_{\alpha}([f_h(\mathbf{x}^*)]_{h \in \mathcal{H}'})$ .

---

### 3.4. Enhancements

As the vanilla scenario decomposition algorithm (Algorithm 1) is quite slow, we introduce several enhancements to improve its computational efficiency in this section.

**3.4.1. Cartesian Cuts** By construction, each no-good cut (5) eliminates exactly one solution in Algorithm 1. The main idea behind the Cartesian Cuts is to determine identify “similar” solutions and eliminate them at the same time using a single inequality.

To be more precise, suppose that we solve scenario batch subproblems S-RA $_{\alpha}(\mathcal{H}^p)$  each  $p = 1, \dots, P$  and obtain a set of solutions  $\mathcal{T} := \{\mathbf{x}^{(p)} : p = 1, \dots, P\}$ . Then, using the K-means algorithm,

we cluster these solutions into the clusters  $\mathcal{T}_c$  for some  $c \in \mathcal{C}$ . For each cluster  $c \in \mathcal{C}$ , we obtain the sets  $\mathcal{A}_{c,n} = \{k : x_{n,k} = 1 \text{ for some } \mathbf{x} \in \mathcal{T}_c\}$  and  $\mathcal{K}_c = \mathcal{A}_{c,1} \times \cdots \times \mathcal{A}_{c,N}$ . Then, we add the following no-good cut:

$$\sum_{n=1}^N \sum_{k \in \mathcal{A}_{c,n}} x_{n,k} \leq N-1 \quad c \in \mathcal{C}. \quad (6)$$

However, due to the Cartesian product operation, this inequality may cut solutions that do not come from the optimal solutions of scenario subproblems. Therefore, it is not valid unless the objective values of the solutions in the sets  $\mathcal{K}_c \setminus \mathcal{T}_c$  are computed and LB is updated, if needed.

To summarize, the Cartesian Cuts (6) have the advantage of eliminating more solutions per inequality (and, hence, per iteration) compared to the standard no-good cut (5). Therefore, they may have a positive effect on both LB and UB progress in Algorithm 1, and have the potential of reducing the number of iterations needed for convergence. On the other hand, the implementation of Cartesian Cuts has an overhead due to the computation of  $s_\alpha([f_h(\mathbf{x})]_{h \in \mathcal{H}})$  for every  $\mathbf{x} \in \mathcal{K}_c \setminus \mathcal{T}_c$ .

### 3.4.2. Symmetry Breaking Inequalities and Symmetry-Enhanced Cartesian Cuts

Interestingly, “no in-take” action may be a part of an optimal solution for the Antibiotics Time Machine Problem (see, e.g. Figure 5a in Mesüm et al. (2024)), meaning that we may opt for a shorter treatment plan than the one allowed with  $N$ . This action can be modeled by adding a  $d \times d$  identity matrix, denoted by  $I$ , to the set of antibiotic matrices (we will refer to the index of this matrix by  $I$ ). However, since the identity matrix is commutative with every matrix (e.g., applying antibiotics in the order  $1-I-1$  or  $1-1-I$  have the same effect), this results in symmetry in the problem, a known issue for the MILP solvers. In order to resolve this issue, we add the following symmetry breaking constraints:

$$x_{1,I} = 0 \text{ and } x_{n,I} \leq x_{n+1,I} \quad n = 1, \dots, N-1. \quad (7)$$

Once inequalities (7) are considered, it is possible to further strengthen the Cartesian cuts (6) introduced before as given in the following proposition:

**PROPOSITION 3.** *Consider a cluster  $\mathcal{T}_c$  of solutions as described before. Let  $N^*$  denote the first position that the identity matrix is present in  $\mathcal{T}_c$ , that is,*

$$N^* = \begin{cases} n, & \text{if } I \notin \mathcal{A}_{c,n-1}, I \in \mathcal{A}_{c,n} \\ N/A, & \text{otherwise.} \end{cases}$$

*Let  $N^I$  denote the first position for which the identity matrix is the singleton in the sets  $\mathcal{A}_{c,1}, \dots, \mathcal{A}_{c,N}$ , that is,*

$$N^I = \begin{cases} n, & \text{if } \{I\} \neq \mathcal{A}_{c,n-1}, \{I\} = \mathcal{A}_{c,n}, \\ N/A, & \text{otherwise.} \end{cases}$$

Consider the set  $\mathcal{X}' = \{\mathbf{x} \in \mathcal{X} : (6), (7)\}$  and the following set definitions:

$$\tilde{\mathcal{A}}_{c,n}^{\bar{N}} := \begin{cases} \mathcal{A}_{c,n}^N & \text{if } n \in \{1, N^* - 1, \dots, \bar{N} - 1\} \\ \mathcal{A}_{c,n}^N \cup \{I\} & \text{if } n \in \{2, \dots, N^* - 2\} \\ \{I\} & \text{if } n = \bar{N} \end{cases}, \tilde{\mathcal{A}}_{c,n}^N := \begin{cases} \mathcal{A}_{c,n} & \text{if } n \in \{1, N^* - 1, \dots, N\} \\ \mathcal{A}_{c,n}^N \cup \{I\} & \text{if } n \in \{2, \dots, N^* - 2\} \end{cases}.$$

Then, we have the following:

(i) If  $N^*$  is defined but  $N^I$  is not defined, then the following inequalities are valid for  $\mathcal{X}'$ :

$$\sum_{n=1}^{\bar{N}-1} \sum_{k \in \tilde{\mathcal{A}}_{c,n}^{\bar{N}}} x_{n,k} \leq \bar{N} - 1, \quad \bar{N} = N^*, \dots, N - 1 \quad (8a)$$

$$\sum_{n=1}^N \sum_{k \in \tilde{\mathcal{A}}_{c,n}^N} x_{n,k} \leq N - 1. \quad (8b)$$

(ii) If both  $N^*$  and  $N^I$  are defined, then the following inequalities are valid for  $\mathcal{X}'$ :

$$\sum_{n=1}^{\bar{N}-1} \sum_{k \in \tilde{\mathcal{A}}_{c,n}^{\bar{N}}} x_{n,k} \leq \bar{N} - 1, \quad \bar{N} = N^*, \dots, N^I. \quad (9)$$

PROOF. We will first prove the validity of inequality (8a) by showing that it is valid for both  $\mathcal{X}'_0 := \mathcal{X}' \cap \{x_{\bar{N},I} = 0\}$  and  $\mathcal{X}'_1 := \mathcal{X}' \cap \{x_{\bar{N},I} = 1\}$  (notice that we have  $\tilde{\mathcal{A}}_{c,n}^{\bar{N}} = \{I\}$  in this case).

- Case 0: For the set  $\mathcal{X}'_0$ , the validity trivially follows since the remaining inequality  $\sum_{n=1}^{\bar{N}-1} \sum_{k \in \tilde{\mathcal{A}}_{c,n}^{\bar{N}}} x_{n,k} \leq \bar{N} - 1$  is simply redundant.
- Case 1: For the set  $\mathcal{X}'_1$ , consider an element  $\gamma \in \tilde{\mathcal{A}}_{c,n}^1 \times \dots \times \tilde{\mathcal{A}}_{c,n}^{\bar{N}}$ . We have two subcases:
  - Case 1a: We have  $(\gamma_1, \dots, \gamma_{\bar{N}-1}) \in \Gamma := \mathcal{A}_{c,1} \times \dots \times \mathcal{A}_{c,\bar{N}-1}$ . Since  $(\gamma_1, \dots, \gamma_{\bar{N}-1}, I, \dots, I) \in \mathcal{K}_c$ , inequality (8a) is valid as it does not cut off any new solutions.
  - Case 1b: We have  $(\gamma_1, \dots, \gamma_{\bar{N}-1}) \notin \Gamma$ , which implies that there exists  $n' \in \{2, \dots, N^* - 2\}$  such that  $\gamma_{n'} = I$ . Since  $\gamma_{n'} = I$  and  $I \notin \tilde{\mathcal{A}}_{c,N^*-1}^{\bar{N}}$ , inequality (8a) can only cut off solutions that violate the symmetry breaking constraint (7). Therefore, it remains valid.

In both cases, inequality (8a) does not cut off any feasible solutions other than the ones coming from  $\mathcal{K}_c$ . Therefore, it is valid.

The validity of inequality (8b) can be proven by repeating Case 1b of the proof of the validity of inequality (8a). The validity of inequality (9) can be proven by repeating the proof of the validity of inequality (8a) by replacing  $N - 1$  with  $N^I$ .  $\square$

Note that inequalities (8)-(9) are stronger than the Cartesian cuts (6).

**3.4.3. Scenario Regrouping** An interesting property of Algorithm 1 is that the upper bound obtained via Proposition 2 depends on the composition of the batches. In particular, even under the same solution, by regrouping batches, we may obtain different upper bounds. For example,

suppose that we have eight scenarios and two batches with the same optimal solutions that give the values  $s_{0.25}([0.9, 1.0, 1.0, 1.0]) = 0.9$  and  $s_{0.25}([0.5, 0.5, 0.8, 0.8]) = 0.5$ , which lead to the upper bound  $UB = \frac{0.9+0.5}{2} = 0.7$ . Notice that these batches are *unbalanced* as the scenarios with higher (resp. lower) probabilities are in Batch 1 (resp. 2). However, if we regroup the scenarios in a *balanced* manner, we obtain the values  $s_{0.25}([0.5, 0.8, 0.9, 1.0]) = s_{0.25}([0.5, 0.8, 1.0, 1.0]) = 0.5$ , which lead to the upper bound  $UB = \frac{0.5+0.5}{2} = 0.5$ .

Motivated by our observation that the optimal solution of one of the batch subproblems from the first iteration is generally optimal for the overall problem, we propose regrouping the scenarios in such a way that each batch is balanced, in the sense that, the scenarios in each batch share a similar distribution. In order to achieve this, we first compute that objective function value of each scenario  $h$  and sort them. Let the sorted scenarios be labeled as  $(h)$ . Then, we regroup the scenarios into batches according to the following rule for  $p = 1, \dots, P$ :  $\mathcal{H}^p = \{(h) : (h) - 1 \pmod P \equiv p - 1\}$ .

**3.4.4. Warm Start** Notice that any feasible solution obtained for the instance with  $N$  transitions can be extended to a feasible solution with the instance with  $N + 1$  transitions by simply appending an identity matrix for the last position. Motivated by this observation, we record the feasible solutions and the cuts generated by Algorithm 1 for the instance with  $N$  transitions, and *warm start* the algorithm for the instance with  $N + 1$  transitions by the lower bound obtained from the solutions and with the addition of the no-good cuts. This enhancement has the potential to help the progress on both lower and upper bounds, hence, decrease the number of iterations, at the cost of solving bulkier MILP models.

## 4. Dynamic Stochastic Problem

The static version of the problem studied in the previous section determines an antibiotic sequence at the beginning of the treatment. In the dynamic version studied in this section, it is allowed to observe the genotype of the bacteria after each antibiotic application. We present how to formulate and solve the dynamic risk-averse version in Sections 4.1. Then, similar to the static version, we propose a scenario decomposition algorithm in Section 4.2 and several algorithmic enhancements to increase its computational efficiency in Section 4.3.

### 4.1. Risk-Averse Version

The risk-averse dynamic formulation of the Antibiotics Time Machine Problem can be stated similar to problem (4) as

$$\max_{\mathbf{y} \in \mathcal{Y}} s_\alpha([\ell_h(\mathbf{y})]_{h \in \mathcal{H}}),$$

where the set  $\mathcal{Y} := \{\mathbf{y} \in \{0, 1\}^{K \times d} : \sum_{k=1}^K y_{k,i} = 1, i = 1, \dots, d\}$  ensures that exactly one antibiotic  $k$  is chosen for each genotype  $i$  with the introduction of the binary variable  $y_{k,i}$ . Here,  $\ell_h(\mathbf{y})$  can be

computed as  $u_{N,i_w}^h$ , where  $i_w$  is the ID of the wild type, by fixing the  $\mathbf{y}$  decisions and the initial condition  $\mathbf{u}_n^h = \mathbf{r}$  in the following recursion ( $\ell_h(\mathbf{y})$  can also be computed as the output of a forward DP, similar to the one used in Mesüm et al. (2024)):

$$u_{n,j}^h = \sum_{k=1}^K \sum_{i=1}^d u_{n-1,i}^h M_{ij}^{k,h} y_{k,i} \quad n = 1, \dots, N, j = 1, \dots, d, h \in \mathcal{H}. \quad (10)$$

We now model the dynamic version of problem as an MILP. Let  $\mathbf{u}_n$  denote the probability distribution after the  $n$ -th antibiotic is applied as in the static version. To linearize the recursion (10), we define the decision variables  $w_{n-1,k,i}^h$  and utilize the equation:  $w_{n-1,k,i}^h = u_{n-1,i}^h y_{k,i}$ . We then provide an MILP formulation for the risk-averse dynamic version of the problem as below:

$$\text{D-RA}_\alpha(\mathcal{H}) : \max_{\mathbf{u}, \mathbf{w}, \mathbf{y} \in \mathcal{Y}, \lambda, \mu} \lambda - \frac{1}{\alpha |\mathcal{H}|} \sum_{h \in \mathcal{H}} \mu_h \quad (11a)$$

$$\text{s.t. } \lambda - \mu_h \leq \mathbf{q}^\top \mathbf{u}_N^h \quad h \in \mathcal{H} \quad (11b)$$

$$\mu_h \geq 0 \quad h \in \mathcal{H} \quad (11c)$$

$$\mathbf{u}_0^h = \mathbf{r} \quad h \in \mathcal{H} \quad (11d)$$

$$u_{n,j}^h = \sum_{k=1}^K \sum_{i=1}^d M_{ij}^{k,h} w_{n-1,k,i}^h \quad n = 1, \dots, N, j = 1, \dots, d, h \in \mathcal{H} \quad (11e)$$

$$w_{n-1,k,i}^h \geq 0, \quad w_{n-1,k,i}^h \leq y_{k,i} \quad n = 1, \dots, N, k = 1, \dots, K, i = 1, \dots, d, h \in \mathcal{H} \quad (11f)$$

$$\sum_{k=1}^K w_{n-1,k,i}^h = u_{n-1,i}^h \quad n = 1, \dots, N, i = 1, \dots, d, h \in \mathcal{H} \quad (11g)$$

$$\sum_{j=1}^d u_{n,j}^h = 1 \quad n = 1, \dots, N, h \in \mathcal{H} \quad (11h)$$

$$\mathbf{u}_n^h \in \mathbb{R}_+^d \quad n = 1, \dots, N, h \in \mathcal{H}. \quad (11i)$$

In this formulation, the objective function (11a) together with constraints (11b)-(11c) model the risk-averse objective function. Constraint (11d) ensures that the initial state distribution is  $\mathbf{r}$ . Constraints (11d)-(11g) linearize the recursion (10). Constraint (11h) is a valid equality which enforces that  $\mathbf{u}_n^h$  is a probability vector. Finally, constraint (11i) gives the variable domain restrictions.

Although the risk-neutral dynamic version can be modeled and solved as a Markov Decision Process as in Mesüm et al. (2024), the DP proposed in that paper does not directly extend to the risk-averse objective we study in this paper. Therefore, we will resort to the MILP proposed above and propose a scenario batch decomposition algorithm to solve it efficiently.

## 4.2. Scenario Decomposition Algorithm

As in the static version,  $\text{D-RA}_\alpha(\mathcal{H})$  is also challenging to solve. We once again propose a scenario decomposition algorithm that relies on lower and upper bounds derived in the following proposition (we omit the proof of this proposition due to its similarity to that of Proposition 2):

PROPOSITION 4. Let  $\alpha \in (0, 1]$ . Denote the optimal value of problem  $D\text{-RA}_\alpha(\mathcal{H})$  as  $\hat{z}_{\mathcal{H}}$  and  $\mathbf{y}^{(p)} = \arg \max_{\mathbf{y} \in \mathcal{Y}} s_\alpha([\ell_h(\mathbf{y})]_{h \in \mathcal{H}^p})$ . Then,

$$\underbrace{\max_{p=1, \dots, P} \{s_\alpha([\ell_h(\mathbf{y}^{(p)})]_{h \in \mathcal{H}})\}}_{LB} \leq \hat{z}_{\mathcal{H}} \leq \underbrace{\frac{1}{P} \sum_{p=1}^P s_\alpha([\ell_h(\mathbf{y}^{(p)})]_{h \in \mathcal{H}^p})}_{UB}.$$

Our dynamic scenario decomposition algorithm is given below:

---

**Algorithm 2** Dynamic Scenario Decomposition Algorithm.

---

Set  $t = 1$ ,  $\mathbf{y}^* = 0$ ,  $LB = 0$  and  $UB = 1$ .

**while**  $UB - LB > \epsilon$  and  $t \leq \tau$  **do**

Solve  $D\text{-RA}_\alpha(\mathcal{H}^p)$  in parallel for each  $p \in P$ .

If needed, update  $LB$  and  $UB$  as in Proposition 4 and the best feasible solution  $\mathbf{y}^*$

Add the following no-good cut to the set  $\mathcal{Y}$  for each distinct solution  $\mathbf{y}^{(p)}$  obtained:

$$\sum_{(k,i): y_{k,i}^{(p)}=1} y_{k,i} \leq d - 1. \quad (12)$$

Increment  $t = t + 1$ .

Carry out an “out-of-sample” analysis with a new sample  $\mathcal{H}'$  by computing  $s_\alpha([\ell_h(\mathbf{y}^*)]_{h \in \mathcal{H}'})$ .

---

### 4.3. Enhancements

For the dynamic version of the problem, two of the enhancements introduced for the static version are directly applicable: i) scenario regrouping (Section 3.4.3), ii) warm start (Section 3.4.4). We would like to note that for the latter enhancement, unlike the static version, we cut the solutions coming from smaller instances after evaluating their objective function values since these values may improve as  $N$  is increased. Although the Cartesian cut enhancement (Section 3.4.1) is also applicable for the dynamic version in principle, our preliminary experiments have showed that the number of solutions in clusters increases rapidly (as  $d > N$  in our computational setting), resulting in prohibitively large number evaluations, and, therefore, slowing down our algorithm.

As an additional enhancement, we introduce the notion of *irrelevant* genotypes, which we now explain. Given an initial genotype  $\mathbf{i}_I \in \{0, 1\}^a$  and the length of treatment  $N$ , we define the set of irrelevant genotypes as  $\mathcal{I}_N(\mathbf{i}_I) := \{\mathbf{g} \in \{0, 1\}^a : \|\mathbf{i}_I - \mathbf{g}\|_1 + \|\mathbf{g}\|_1 > N\}$ . Notice that these are exactly the genotypes for which the number of transitions from  $\mathbf{i}_I$  to  $\mathbf{g}$  and then to  $\mathbf{0}$  requires more than  $N$  antibiotic applications. Since this happens with zero probability no matter what the policy is, we can safely set  $y_{I,i} = 1$  from the very beginning for the ID of such genotypes. To give an example,

for the instance in Figure 1, if the initial genotype is 001 and  $N = 2$ , then the genotype 101 is irrelevant. The identification of irrelevant genotypes is crucial for the progress of the upper bound in Algorithm 2, since, otherwise, the algorithm might stall.

## 5. Computational Results

In this section, we provide the results of our computational experiments. We use a 64-bit workstation with two Intel(R) Xeon(R) Gold 6248R CPU (3.00GHz) processors (256 GB RAM) and the Python programming language. We utilize Gurobi 11 to solve the MILPs with the default settings, except the absolute optimality gap parameter is set to 0.001 and the time limit is set to two hours. We set  $|\mathcal{H}| = |\mathcal{H}'| = 2000$ ,  $|\mathcal{H}^p| = 50$ ,  $\alpha = 0.1$ ,  $\epsilon = 0.01$  and  $\tau = 5$ . We use the experimental data from Mira et al. (2017) with  $K = 23$  antibiotics and  $d = 2^4$  genotypes, and sample the transition probability matrices as explained in Appendix A. We also add the identity matrix as an additional antibiotic to model the “no in-take” action.

### 5.1. Static Version

In this section, we provide the results of the computational experiments that we have conducted for the static version of the problem.

**5.1.1. Computational Effort** In this part, we first analyze the effect of enhancements we introduce in Section 3.4. For this purpose, we run our algorithm under the following settings:

- Setting 0: No enhancements
- Setting 1: All enhancements except Cartesian Cuts (Section 3.4.1)
- Setting 2: All enhancements except symmetry breaking inequalities and symmetry-enhanced Cartesian Cuts (Section 3.4.2)
- Setting 3: All enhancements except scenario regrouping (Section 3.4.3)
- Setting 4: All enhancements except warm start (Section 3.4.4)
- Setting 5: All enhancements

In Table 1, for two initial genotype  $i_I$  and treatment length  $N$  combinations, we report the number of iterations needed for Algorithm 1 to terminate (in column “#”) along with the CPU time in seconds (in column “Time”).

**Table 1** Effect of enhancements for the risk-averse static version in terms of computational effort.

$N$	$i_I$	Setting 0		Setting 1		Setting 2		Setting 3		Setting 4		Setting 5	
		#	Time	#	Time	#	Time	#	Time	#	Time	#	Time
4	0001	4	2441	1	593	2	1246	1	694	2	1153	1	697
5	1111	5	5063	4	3538	5	5482	4	3676	3	2754	3	2650

As can be seen from Table 1, Setting 5 is significantly more efficient than Setting 0 for the first instance. In terms of individual enhancements, we observe that Setting 2 and Setting 4 require more iterations than Setting 5, indicating the importance of symmetry exploitation and warm start, respectively. We also note that Setting 1 is faster than Setting 5, suggesting that the overhead of implementing the Cartesian cuts exceeds their benefit in this case. As for the second instance, we observe that the exclusion of any enhancement increases the solution time, suggesting that the enhancements might be even more influential for instances with larger values of  $N$ . Unlike the first instance, we see that Setting 1 and Setting 3 are also slower than Setting 5 for the second instance, indicating the importance of the Cartesian cuts and scenario regrouping. We note that, in addition to these two instances, we report the results of the computational experiments for each of the settings introduced above for each initial genotype with  $N = 4$  in Appendix B as well.

Next, we run the risk-averse model for  $N = 5, \dots, 8$  using all the enhancements (i.e., Setting 5), and report the results in Table 2. We report the number of iterations, or the difference between UB and LB (in *italic*) after the last iteration if this difference is larger than  $\epsilon$ , under the column “#/%”. Despite the computational benefits brought by the enhancements, it is still challenging to solve the instances with larger  $N$  values. Therefore, we apply the following filtering rules to eliminate the antibiotics that do not appear in the optimal solutions of the instances with smaller  $N$  values:

- Filter I: This filter is applied to instances with  $N = 5, 6$  and it is *genotype-independent*, that is, the same filtered set of antibiotics is used for any initial genotype  $\mathbf{i}_I$ . In particular, we restrict the set of antibiotics to those that appear in the optimal solutions of instances with  $N = 4$  for any  $\mathbf{i}_I$ . This reduces the number of antibiotics from 23 to 13.
- Filter II: This filter is applied to instances with  $N = 7, 8$  and it is *genotype-dependent*, that is, a different filtered set of antibiotics is used for each initial genotype  $\mathbf{i}_I$ . In particular, we restrict the set of antibiotics for an initial genotype  $\mathbf{i}_I$  to those that appear in the optimal solutions of instances with  $N = 4, 5, 6$  for the same initial genotype  $\mathbf{i}_I$ . This reduces the number of antibiotics to a range of one to six, depending on the initial genotype.



**Table 2** Computational effort to solve the risk-averse static version for  $N = 5, \dots, 8$  with different filtering mechanisms (no iteration limit is enforced for  $N = 5$ ).

$i_I$	$N = 5$		$N = 5$ - Filter I		$N = 6$ - Filter I		$N = 7$ - Filter II		$N = 8$ - Filter II	
	#/%	Time	#/%	Time	#/%	Time	#/%	Time	#/%	Time
1000	1	5360	1	1045	1	6211	1	4	1	4
0100	1	2077	1	442	1	2605	1	4	1	4
0010	1	1647	1	408	1	2143	1	103	1	239
0001	1	3228	1	615	1	3835	2	1065	2	4184
1100	1	3705	1	715	1	4189	1	9	1	30
1010	1	1985	1	387	1	2633	2	453	2	1168
1001	2	7586	2	1667	2	9523	2	5585	0.15	36398
0110	1	2213	1	378	1	1730	1	49	1	76
0101	3	14816	1	1033	1	3856	1	1484	2	9796
0011	2	4296	2	940	2	5380	2	449	2	1345
1110	2	5689	2	1005	2	6582	2	607	2	1952
1101	3	14438	2	1794	4	21925	4	3410	5	14021
1011	1	2004	1	406	2	5312	2	584	2	1546
0111	2	4780	2	850	2	5572	2	999	3	6231
1111	9	47852	3	2650	3	16716	3	3588	0.02	26115
<b>Avg</b>	2.1	8112	1.5	956	1.7	6548	1.8	1226	2.3	6874

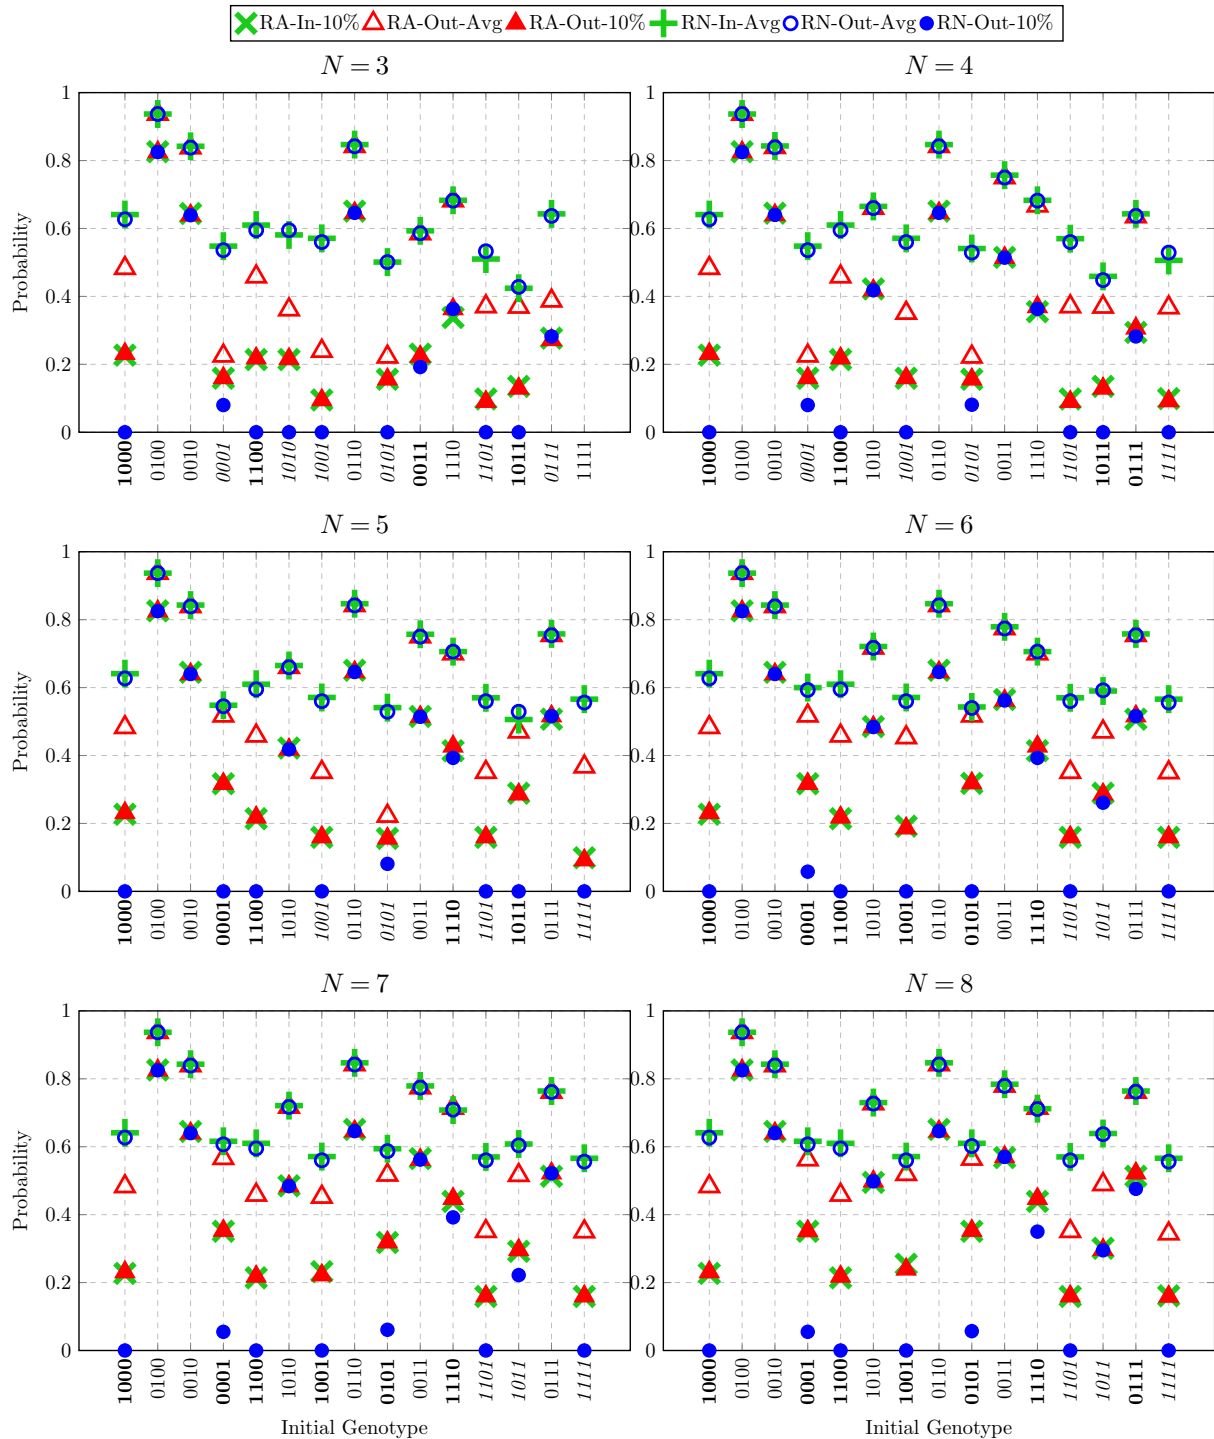
Although we introduce the filtering mechanisms to relieve the computational burden, there is also a biological reasoning: some of the antibiotics from the data set of Mira et al. (2017) have the same active ingredient, therefore, it is conceivable that only some of those would appear in the optimal solutions. A further justification of the filtering mechanism can be found in the first two columns of Table 2, where we compare the computational effort for  $N = 5$  without and with filtering. In fact, the solutions obtained under these two experiment turn out to be identical (recall that we have decided the filtering based on the optimal solutions of instances with  $N = 4$ ), yet the CPU time is decreased from 8112 seconds to 956 seconds on average. We note that filtering mechanisms allow us to solve the corresponding instances within a reasonable time. The only two time-outs are observed when  $N = 8$  and the initial genotypes are 1001 and 1111.

**5.1.2. Comparison of Risk-Averse and Risk-Neutral Solutions** In this section, we compare risk-averse and risk-neutral solutions from different perspectives. Our results are visualized in Figure 3, where we report six values for each initial genotype and  $N = 3, \dots, 8$ :

- RA-In-10%: The optimal value of the “in-sample” optimization problem under the risk-averse objective.
- RA-Out-Avg: The average value of all “out-of-sample” scenarios when the risk-averse solution is applied.
- RA-Out-10%: The average value of the worst 10% “out-of-sample” scenarios when the risk-averse solution is applied.
- RN-In-Avg: The optimal value of the “in-sample” optimization problem under the risk-neutral objective.

- RN-Out-Avg: The average value of all “out-of-sample” scenarios when the risk-neutral solution is applied.
- RN-Out-10%: The average value of the worst 10% “out-of-sample” scenarios when the risk-neutral solution is applied.

**Figure 3** Comparison of risk-averse and risk-neutral solutions for the static version.



We have several observations from our analysis: Firstly, our approach is quite reliable as the in-sample and out-of-sample performances under both risk-averse and risk-neutral objectives are close to each other. In fact, RA-In-10% values, symbolized with a cross sign, and RA-Out-10% values, symbolized with a filled triangle, almost always coincide (similarly, RN-In-Avg values, symbolized with a plus sign, and RN-Out-Avg values, symbolized with an empty circle, coincide as well). Therefore, our results suggest that the choice of the scenario size as  $|\mathcal{H}| = 2000$  is reasonable.

Secondly, the results illustrate an expected trade-off between risk-averse and risk-neutral objectives. Clearly, the risk-averse solution has a better RA-Out-10% value compared to the RN-Out-10% value obtained from the risk-neutral solution. However, this comes with a sacrifice from the average performance as RN-Out-Avg values are typically larger than RA-Out-Avg values. Depending on how this trade-off plays out, we roughly divide the genotypes into three groups:

- Group 1: These are *indifferent* genotypes for which risk-averse and risk-neutral solutions are indistinguishable (e.g., genotypes 0100, 0010, 0110, 1110 for  $N = 3$ ).
- Group 2: These are *good* genotypes for which the increase in the worst 10% performance is higher than the decrease in the average performance when switched from the risk-neutral solution to the risk-averse solution (e.g., genotypes 1000, 1100, 0011, 1011 for  $N = 3$ ). The labels of these genotypes are boldface.
- Group 3: These are *bad* genotypes for which the increase in the worst 10% performance is lower than the decrease in the average performance when switched from the risk-neutral solution to the risk-averse solution (e.g., genotypes 0001, 1010, 1001, 0101, 1101, 0111 for  $N = 3$ ). The labels of these genotypes are italicized.

We notice that the number of *bad* (resp. *good*) genotypes decreases (resp. increases) consistently as  $N$  increases. For  $N = 4$ , the number of *good* and *bad* genotypes are four and five, respectively. Compared to the risk-neutral approach, solving the problem in a risk-averse manner improves the worst 10% performance by 7.3%, computed as the difference of RA-Out-10% and RN-Out-10%, on average whereas it decreases the average performance by 10.3%, computed as the difference of RA-Out-Avg and RN-Out-Avg. Going up to  $N = 6$ , the number of *good* and *bad* genotypes become six and three, respectively. Compared to the risk-neutral solution, the risk-averse solution improves the worst 10% performance by 10.6% on average while decreasing the average performance by 6.8%. Finally, solving the problem for  $N = 8$ , we see that the number of *good* genotypes increases to seven while the number of *bad* genotypes is still three. The risk-averse solution is better than the risk-neutral solution by 11.6% in terms of the worst 10% performance, which comes with a 6.5% sacrifice from the average performance.

As illustrated above, switching from the risk-neutral solution to the risk-averse solution comes with a significant gain from the worst 10% performance albeit a considerable sacrifice from the

average performance in most cases even for *good* genotypes. These observations illustrate the price of risk-aversion in the static version. However, as  $N$  increases, the aforementioned gain increases and the sacrifice decreases on average.

We remark that we do not solve the problem for larger  $N$  values since it would require too much computational effort whereas the potential benefit is projected to be small (notice that the objective functions values seem to stabilize for the larger  $N$  values reported in Figure 3).

## 5.2. Dynamic Version

In this section, we present the results for the dynamic version of the problem. We note that as this approach allows for intermediate observations, we can take more informed steps, and, therefore, the success of the dynamic version will be higher compared to the success of the static version in terms of reversing the antibiotic resistance.

**5.2.1. Computational Effort** We run Algorithm 2 with the enhancements introduced in Section 4.3 for  $N = 3, \dots, 8$ , and report the results in Table 3, whose format is similar to that of Table 2.

**Table 3** Computational effort to solve the risk-averse dynamic version for  $N = 3, \dots, 8$  with different filtering mechanisms.

$i_I$	$N = 3$		$N = 4$		$N = 5 - \mathbf{F}$		$N = 6 - \mathbf{F}$		$N = 7 - \mathbf{F}$		$N = 8 - \mathbf{F}$	
	#/%	Time	#/%	Time	#/%	Time	#/%	Time	#/%	Time	#/%	Time
<b>1000</b>	0.02	466	0.02	613	0.01	640	0.02	819	0.03	1122	0.03	1181
<b>0100</b>	0.02	923	0.02	1459	1	121	1	158	1	211	1	236
<b>0010</b>	2	184	2	231	1	120	1	148	1	180	1	204
<b>0001</b>	0.03	914	0.03	1368	0.01	656	4	668	0.01	1115	0.01	1161
<b>1100</b>	3	241	0.04	5380	0.03	642	0.04	979	0.04	1216	0.04	1442
<b>1010</b>	1	83	0.03	9369	1	125	0.02	1032	0.02	1285	0.02	1564
<b>1001</b>	2	173	0.03	18593	0.03	645	0.03	932	0.03	1161	0.03	1333
<b>0110</b>	1	80	0.01	1999	1	119	1	162	1	190	1	234
<b>0101</b>	2	171	0.05	15262	0.02	634	0.03	1008	0.03	1232	2	547
<b>0011</b>	2	165	0.03	1877	1	123	0.02	848	0.02	1018	0.02	1103
<b>1110</b>	0.03	812	0.03	1406	0.01	724	0.02	1114	0.05	1569	0.05	1845
<b>1101</b>	0.05	840	0.05	1690	0.03	668	0.06	993	0.05	1323	0.05	1439
<b>1011</b>	0.03	821	0.03	1233	0.02	728	0.02	904	0.02	1246	0.02	1287
<b>0111</b>	4	589	4	812	1	166	0.03	937	0.04	1342	0.04	1436
<b>1111</b>	-	-	0.06	53206	0.02	787	0.07	1105	0.07	1381	0.06	1600
<b>Avg</b>	3.4	461	4.7	7633	3.4	459	4.1	787	4.2	1039	4.0	1107

The dynamic version turns out to be more challenging than the static version due to slower progress on the upper bound, therefore, we frequently reach the iteration limit of  $\tau$  even for small values of  $N = 3, 4$ . In addition, solving the subproblems becomes time consuming for larger values of  $N$  as well, therefore, we introduce another filtering mechanism: We first solve the risk-averse dynamic version for  $N = 4$  and record the solutions as  $\mathbf{y}^{[i_I]}$  when the initial genotype is  $i_I$ . Then, if

an antibiotic  $k$  is used when the stochastic process is at genotype  $i$  for *any* initial genotype, we limit the set of available antibiotics for genotype  $i$  to be  $K_i := \{k = 1, \dots, K : y_{k,i}^{[i_I]} = 1 \text{ for some } i_I\}$  except the genotype 1111, for which we do not apply any filter. This filtering rule limits the antibiotics to be used and decreases the solution time of the MILP significantly. The results reported in Table 3 show that although Algorithm 2 reaches the iteration limit in most cases, the gap obtained upon termination is generally small, ranging between 0.01-0.07.

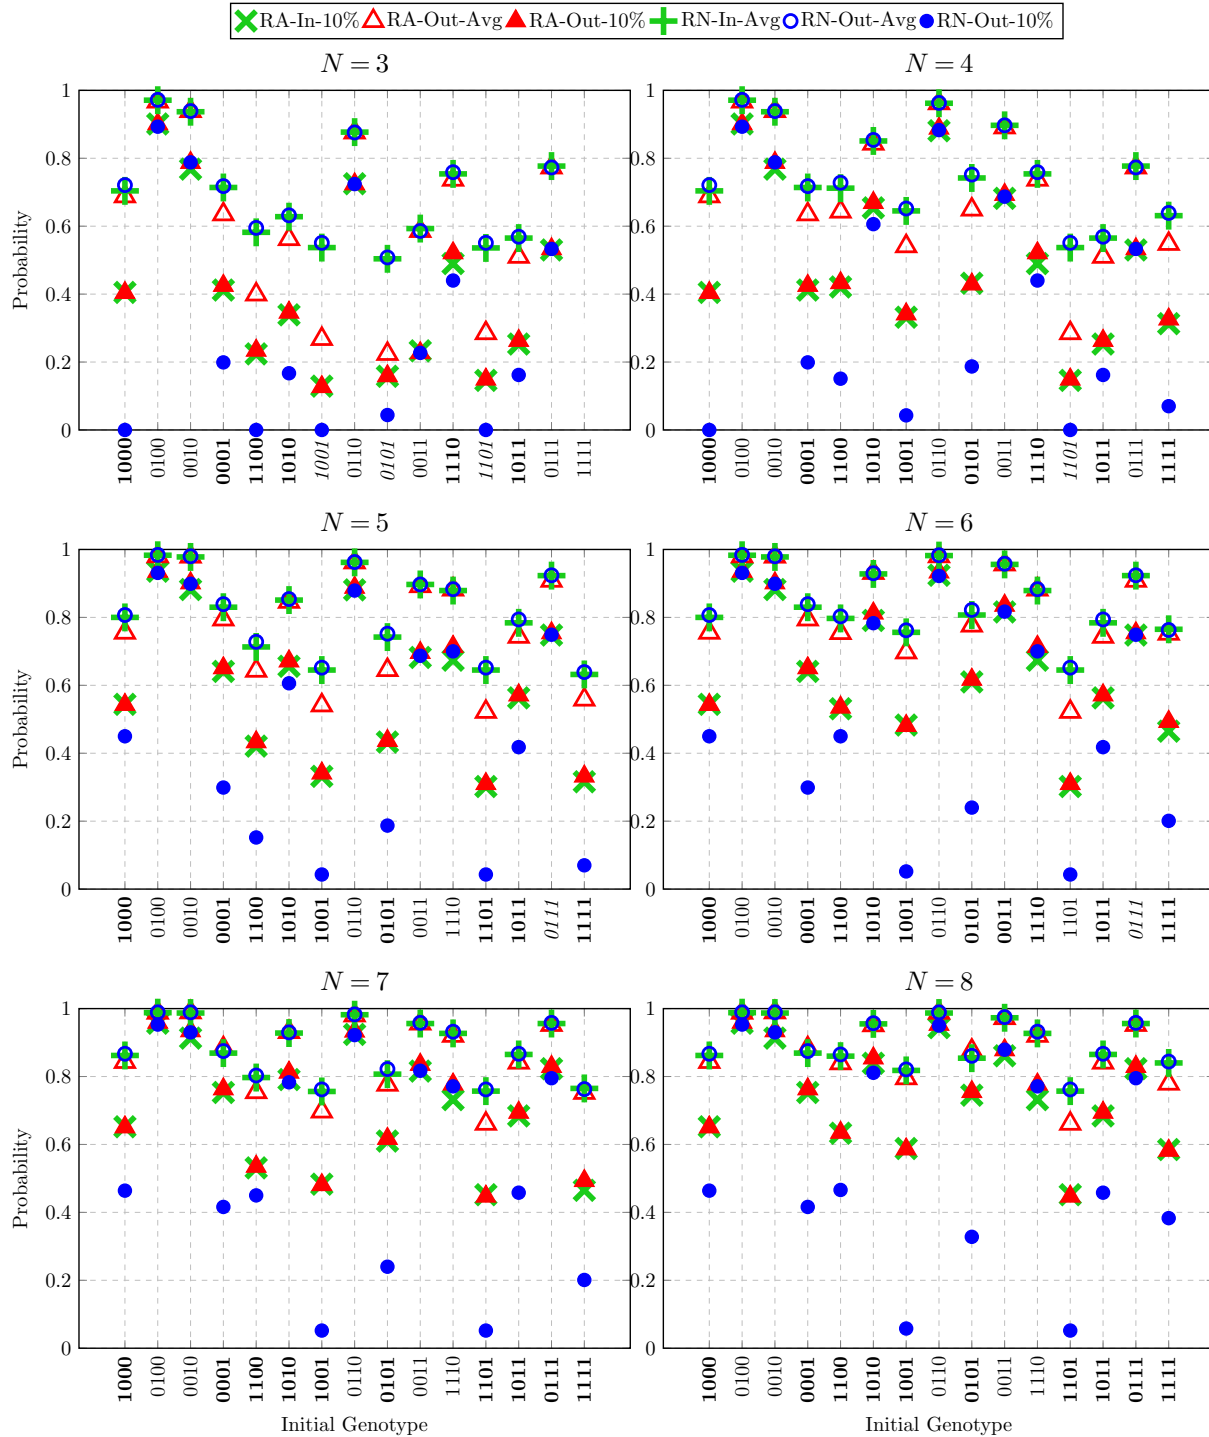
**5.2.2. Comparison of Risk-Averse and Risk-Neutral Solutions** Similar to Figure 3, we now compare risk-averse and risk-neutral solutions for the dynamic version in Figure 4, for each initial genotype and  $N = 3, \dots, 8$ . A significant difference between the results of static and dynamic versions is that both average and worst 10% performances are much higher when intermediate observations are allowed.

As before, the results suggest that our approach is again reliable as the in-sample and out-of-sample performances under both risk-averse and risk-neutral objectives are close to each other. Next, we will discuss the trade-off between risk-averse and risk-neutral objectives. Using the same terminology, we observe that the number of *good* genotypes increases and the number of *bad* genotypes decreases rapidly as  $N$  increases.

For  $N = 4$ , we observe that nine genotypes are *good* and only one genotype is *bad*. Compared to the risk-neutral solution, the risk-averse solution improves the worst 10% performance by 14.0% while losing 5.9% from the average performance. Moving up to  $N = 6$ , the number of *good* genotypes increases to ten while the number of *bad* genotypes stays the same. The risk-averse solution performs 14.2% better in the worst 10% performance with only 3.1% sacrifice from the average performance. Finally, at  $N = 8$ , we have ten *good* and zero *bad* genotypes. The risk-averse solution has a 17.2% better worst 10% performance compared to the risk-neutral solution with only 1.8% loss from the average performance. Therefore, we can say that as  $N$  increases, the benefit of solving the problem in our proposed risk-averse approach increases considerably.

To summarize, when switched from the risk-neutral solution to the risk-averse solution, we observe significantly better worst 10% performance with relatively small reductions from the average performance. These observations suggest that the price of risk-aversion is quite low in the dynamic version, and our proposed approach is arguably more compatible with this version.

Finally, we compare the results of static and dynamic versions in terms of their compatibility with risk-aversion. We observe that risk-averse solutions are somewhat “more successful” with the dynamic version as the number of *bad* genotypes is much smaller and the sacrifice from the average performance becomes almost negligible compared to the gain from the worst 10% performance as  $N$  increases. This suggests that risk-averse objective works much better with the dynamic version.

**Figure 4** Comparison of risk-averse and risk-neutral solutions for the dynamic version.

## 6. Conclusions

In this paper, we study the risk-averse Antibiotics Time Machine Problem. We consider two versions of the problem: static and dynamic. For both versions, we propose scenario-based MILP formulations. Since these formulations do not scale well with scenario size, we develop scenario

batch decomposition algorithms with risk-averse subproblems. We introduce several algorithmic enhancements to further improve the computational efficiency. We test the efficacy of our approach on a real dataset. Our extensive computational experiments suggest that risk-averse solutions provide a significantly better worst-case performance compared to risk-neutral solutions, which comes with a deterioration from the average performance. We observe that the gains in terms of worst-case performance dominate the losses from the average performance as  $N$  increase, and the benefits are even more significant for the dynamic version.

**Funding** This work was supported by the Scientific and Technological Research Council of Turkey [grant number 120C151].

## References

- Aggarwal M, Patra A, Awasthi I, George A, Gagneja S, Gupta V, Capalash N, Sharma P (2024) Drug repurposing against antibiotic resistant bacterial pathogens. *European Journal of Medicinal Chemistry* 279:116833, ISSN 0223-5234.
- Ahmed S (2013) A scenario decomposition algorithm for 0–1 stochastic programs. *Operations Research Letters* 41(6):565–569, ISSN 0167-6377.
- Aulin LBS, Liakopoulos A, van der Graaf PH, Rozen DE, van Hasselt JGC (2021) Design principles of collateral sensitivity-based dosing strategies. *Nat Commun* 12(1):5691.
- Ben-Tal A, Nemirovski A (2001) *Lectures on Modern Convex Optimization* (Society for Industrial and Applied Mathematics).
- Bergstrom CT, Lo M, Lipsitch M (2004) Ecological theory suggests that antimicrobial cycling will not reduce antimicrobial resistance in hospitals. *Proceedings of the National Academy of Sciences* 101(36):13285–13290.
- Brown D (2015) Antibiotic resistance breakers: can repurposed drugs fill the antibiotic discovery void? *Nature Reviews Drug Discovery* 14(12):821–832, ISSN 1474-1784.
- Chang CY, Schiano TD (2007) Review article: drug hepatotoxicity. *Aliment Pharmacol Ther* 25(10):1135–1151.
- Deng Y, Ahmed S, Shen S (2018) Parallel scenario decomposition of risk-averse 0-1 stochastic programs. *INFORMS Journal on Computing* 30(1):90–105, URL <http://dx.doi.org/10.1287/ijoc.2017.0767>.
- Ding Y, Chen J, Wu Q, Fang B, Ji W, Li X, Yu C, Wang X, Cheng X, Yu HD, Hu Z, Uvdal K, Li P, Li L, Huang W (2024) Artificial intelligence-assisted point-of-care testing system for ultrafast and quantitative detection of drug-resistant bacteria. *SmartMat* 5(3):e1214.
- Gillespie JH (1983) A simple stochastic gene substitution model. *Theoretical population biology* 23(2):202–215.

- Gillespie JH (1984) Molecular evolution over the mutational landscape. *Evolution* 38(5):1116–1129.
- Imamovic L, Sommer MO (2013) Use of collateral sensitivity networks to design drug cycling protocols that avoid resistance development. *Science translational medicine* 5(204):204ra132–204ra132.
- Kocuk B (2022) Optimization problems involving matrix multiplication with applications in materials science and biology. *Engineering Optimization* 54(5):786–804.
- Liu G, Catacutan DB, et al. (2023) Deep learning-guided discovery of an antibiotic targeting *acinetobacter baumannii*. *Nature Chemical Biology* 19(11):1342–1350.
- Liu GY, Yu D, Fan MM, Zhang X, Jin ZY, Tang C, Liu XF (2024) Antimicrobial resistance crisis: could artificial intelligence be the solution? *Military Medical Research* 11(1):7.
- Maltas J, Wood KB (2019) Pervasive and diverse collateral sensitivity profiles inform optimal strategies to limit antibiotic resistance. *PLOS Biology* 17(10):1–34.
- Maltas J, Wood KB (2021) Dynamic collateral sensitivity profiles highlight challenges and opportunities for optimizing antibiotic sequences. *bioRxiv* .
- Mesüm O, Atilgan AR, Kocuk B (2024) A stochastic programming approach to the antibiotics time machine problem. *Mathematical Biosciences* 372:109191, ISSN 0025-5564.
- Mira PM, Barlow M, Meza JC, Hall BG (2017) Statistical package for growth rates made easy. *Molecular biology and evolution* 34(12):3303–3309.
- Mira PM, Crona K, Greene D, Meza JC, Sturmfels B, Barlow M (2015) Rational design of antibiotic treatment plans: a treatment strategy for managing evolution and reversing resistance. *PloS one* 10(5):e0122283.
- Murray CJ, Ikuta KS, Sharara F, Swetschinski L, Aguilar GR, Gray A, Han C, Bisignano C, Rao P, Wool E, et al. (2022) Global burden of bacterial antimicrobial resistance in 2019: a systematic analysis. *The lancet* 399(10325):629–655.
- Nichol D, Jeavons P, Fletcher AG, Bonomo RA, Maini PK, Paul JL, Gatenby RA, Anderson ARA, Scott JG (2015) Steering evolution with sequential therapy to prevent the emergence of bacterial antibiotic resistance. *PLoS computational biology* 11(9):e1004493.
- Talat A, Bashir Y, Khan AU (2022) Repurposing of antibiotics: Sense or non-sense. *Front Pharmacol* 13:833005.
- Tran NM, Yang J (2017) Antibiotics time machines are hard to build. *Notices of the AMS* 64(10):1136–1140.
- Tyers M, Wright GD (2019) Drug combinations: a strategy to extend the life of antibiotics in the 21st century. *Nature Reviews Microbiology* 17(3):141–155, ISSN 1740-1534.
- van Duijn PJ, Verbrugghe W, Jorens PG, Spöhr F, Schedler D, Deja M, Rothbart A, Annane D, Lawrence C, Nguyen Van JC, Misset B, Jereb M, Seme K, Šifrer F, Tomić V, Estevez F, Carneiro J, Harbarth S, Eijkemans MJC, Bonten M, Goossens H, Malhotra-Kumar S, Lammens C, Vila J, Roca I (2018)



The effects of antibiotic cycling and mixing on antibiotic resistance in intensive care units: a cluster-randomised crossover trial. *The Lancet Infectious Diseases* 18(4):401–409, ISSN 1473-3099.

Ventola CL (2015) The antibiotic resistance crisis: part 1: causes and threats. *P T* 40(4):277–283.

Wong F, Zheng EJ, Valeri JA, Donghia NM, Anahtar MN, Omori S, Li A, Cubillos-Ruiz A, Krishnan A, Jin W, Manson AL, Friedrichs J, Helbig R, Hajian B, Fiejtek DK, Wagner FF, Soutter HH, Earl AM, Stokes JM, Renner LD, Collins JJ (2023) Discovery of a structural class of antibiotics with explainable deep learning. *Nature* 626(7997):177–185.

Zagajewski A, Turner P, Feehily C, El Sayyed H, Andersson M, Barrett L, Oakley S, Stracy M, Crook D, Nellåker C, Stoesser N, Kapanidis AN (2023) Deep learning and single-cell phenotyping for rapid antimicrobial susceptibility detection in escherichia coli. *Commun Biol* 6(1):1164.

Zhang C, Fu X, Liu Y, Zhao H, Wang G (2024) Burden of infectious diseases and bacterial antimicrobial resistance in china: a systematic analysis for the global burden of disease study 2019. *The Lancet Regional Health – Western Pacific* 43, ISSN 2666-6065.

## Appendix A: Matrix and Scenario Construction

An important aspect of the Antibiotics Time Machine Problem is the calculation of transition probability matrices for which we follow Mira et al. (2015). Suppose that we are given a vector of growth rate measurements of genotypes for a specific antibiotic, denoted by  $\omega \in \mathbb{R}_+^d$ . Let us denote the set of neighboring genotype pairs by  $J$ , and take a pair  $(j, j')$  from the set  $J$ . If  $\omega_{j'} > \omega_j$ , then the transition probability from genotype  $j$  to  $j'$  is positive, otherwise the probability is zero. In our paper, we use the so-called *correlated probability model*, in which the transition probabilities are calculated as follows (Mira et al. 2015):

$$M_{j,j'}(\omega) := \frac{\max\{0, \omega_{j'} - \omega_j\}}{\sum_{j'':(j'',j) \in J} \max\{0, \omega_{j''} - \omega_j\}}, \quad (j, j') \in J.$$

Now, suppose that the growth rate measurement for antibiotic  $k$ , denoted by  $\omega^k$ , is a random variable and its  $i$ -th entry takes a value from the set  $\{w_i^{k,r} : r = 1, \dots, R\}$  with equal probability (see, e.g., Mira et al. (2017), Mesüm et al. (2024)). Assuming independence between entries, we conclude that  $\omega^k$  takes values from a set  $\{\omega^{k,s} : s \in \mathcal{S}\}$  for some index set  $\mathcal{S} := \{1, \dots, R\}^d$ . To give some perspective, Mira et al. (2017) record the growth rates of  $K = 23$  antibiotics for  $R = 12$  times in which case there are  $d = 2^4$  genotypes. Due to the astronomically large size of the set  $\mathcal{S}$ , it is not directly possible to incorporate all scenarios. Instead, we take a sample  $\mathcal{H} \subseteq \mathcal{S}^K$ , where  $|\mathcal{H}| \ll |\mathcal{S}|^K$ .

## Appendix B: Effect of Enhancements for the Static Version with $N = 4$

In order to compare the effect of enhancements, we solve the risk-averse static version for each initial genotype with  $N = 4$  and report the results in Table 4. Comparing Setting 0 and Setting 5, we see considerable savings in terms of computational time. In these experiments, the symmetry breaking constraints and symmetry-enhanced Cartesian Cuts are the most effective enhancements followed by warm starting.

**Table 4** Effect of enhancements for the risk-averse static version in terms of computational effort for  $N = 4$  (no iteration limit is enforced in this experiment).

$i_I$	Setting 0		Setting 1		Setting 2		Setting 3		Setting 4		Setting 5	
	#	Time	#	Time	#	Time	#	Time	#	Time	#	Time
<b>1000</b>	1	478	1	517	1	540	1	527	1	483	1	535
<b>0100</b>	1	107	1	229	1	128	1	231	1	88	1	235
<b>0010</b>	1	223	1	226	1	220	1	220	1	218	1	235
<b>0001</b>	4	2441	1	593	2	1246	1	694	2	1154	1	697
<b>1100</b>	2	1020	1	491	2	1039	1	547	2	1104	1	553
<b>1010</b>	1	323	1	307	1	299	1	281	1	282	1	274
<b>1001</b>	2	1068	2	1052	2	1121	2	1162	2	1193	2	1218
<b>0110</b>	1	302	1	331	1	321	1	314	1	261	1	318
<b>0101</b>	7	4599	3	1840	3	2187	3	2055	3	2094	3	2206
<b>0011</b>	2	590	2	662	2	590	2	649	2	644	2	625
<b>1110</b>	3	1226	2	768	2	742	2	743	2	684	2	734
<b>1101</b>	6	3156	3	1519	5	2927	3	1675	3	1566	3	1631
<b>1011</b>	3	852	2	570	2	647	2	636	2	619	2	627
<b>0111</b>	3	1251	2	848	2	822	2	804	2	816	2	852
<b>1111</b>	4	2163	3	1692	3	1708	3	1808	3	1771	3	1855
<b>Avg</b>	2.7	1320	1.7	776	2.0	969	1.7	823	1.9	865	1.7	840