ARTICLE

Check for updates

A game theoretic approach reveals that discretizing clinical information can reduce antibiotic misuse

Maya Diamant[1], Shoham Baruch[2,3], Eias Kassem [ID] [4], Khitam Muhsen[2], Dov Samet[1], Moshe Leshno[1] & Uri Obolski [ID] [2,3 ✉]

The overuse of antibiotics is exacerbating the antibiotic resistance crisis. Since this problem is a classic common-goods dilemma, it naturally lends itself to a game-theoretic analysis. Hence, we designed a model wherein physicians weigh whether antibiotics should be prescribed, given that antibiotic usage depletes its future effectiveness. The physicians' decisions rely on the probability of a bacterial infection before definitive laboratory results are available. We show that the physicians' equilibrium decision rule of antibiotic prescription is not socially optimal. However, we prove that discretizing the information provided to physicians can mitigate the gap between their equilibrium decisions and the social optimum of antibiotic prescription. Despite this problem's complexity, the effectiveness of the discretization solely depends on the type of information available to the physician to determine the nature of infection. This is demonstrated on theoretic distributions and a clinical dataset. Our results provide a game-theory based guide for optimal output of current and future decision support systems of antibiotic prescription.

[1] Coller School of Management, Tel Aviv University, Tel Aviv, Israel. [2] Faculty of Medicine, School of Public Health, Tel Aviv University, Tel Aviv, Israel. [3] Faculty of Exact Sciences, Porter School of the Environment and Earth Sciences, Tel Aviv University, Tel Aviv, Israel. [4] Department of Pediatrics, Hillel Yaffe Medical Center, Hadera, Israel. ✉email: uriobols@tauex.tau.ac.il

Resistance to antibiotics has significant public health implications worldwide[1,2]. Antibiotic-resistant infections increase mortality rates, lengthen hospital stays and increase treatment costs[3]. When a patient is suspected of having a bacterial infection, bio-specimens are often collected, cultured, and undergo antimicrobial susceptibility testing. Typically, the results are known within several days. Having the results, a physician can prescribe an appropriate definitive treatment. However, these days might be critical, and therefore an initial empirical antibiotic treatment is given before identification of the causative pathogen[4]. The choice of empirical treatment can be based on either the physician's informal assessment of the likelihood of different types of infection, given the patient's medical symptoms, clinical settings, and the results of immediate diagnostic tests[5,6], or the output of a decision support system[7].

Appropriate empirical treatment, that ex-post matches the culture results, is highly important. Mainly, it reduces the mortality rate due to bacterial infections significantly[8–10]. The obvious adverse effect of unnecessary empiric treatment is excessive use of antibiotics, which drives the emergence of antibiotic-resistant pathogens in the general population[1,11], as well as in the treated patients themselves[12–14], through the evolutionary forces of selection. Nonetheless, patients are often unnecessarily treated with antibiotics (i.e., when no bacterial infection is present)[15].

The problem of resistance emergence has set a challenge to the medical decision-making literature, which traditionally focuses on the well-being of the individual patient. Both empirical and theoretical studies have been conducted to find antibiotic prescription strategies that decelerate the process of resistance emergence and spread[16–22] and analyze its economic cost to society[23,24].

Nearly none of the existing models have analyzed the clinical decision problem of initiating antibiotic treatment from a game-theoretic perspective but instead tried to identify the socially optimal policy—one that considers the benefit of both current and future patients. The underlying assumption in such models is that since antibiotics can only be consumed given the physician's prescription, the recommended policy can and will be implemented, despite the inherent conflict between the social interest and the personal interest of a single patient. However, the assumption of compliance to antibiotic usage guidelines may be unwarranted. Physicians have been documented to ignore recommendations of support systems, both with regards to antibiotic therapy[7,25] and in other settings[26–28]. These deviations may be analyzed as rational strategic behavior, using the framework of game theory.

A game-theoretic approach views a physician as a decision-maker (player) who seeks the treatment policy that will optimize the expected utility of his own patients, given the behavior of the other players. The player is not assumed to consider the utility of patients who are not under his direct responsibility. This assumption can be interpreted in two possible ways: the physician may either act myopically as if each patient himself were the decision-maker, or he may act on behalf of both his present and future patients.

The first interpretation was recently modeled in ref. [29], but the model precluded physicians from considering any long-term implication of antibiotic usage on future patients. Furthermore, this model did not consider a potentially crucial factor underlying the clinical decision of whether to administer an antibiotic treatment: the differential information regarding the probability of bacterial infection in each patient. Hence, the modeled treatment strategies were limited to either treat/not treat everyone, or making a random decision. Additionally, in ref. [30] the authors use an evolutionary game theory approach to model antibiotic prescription behavior. However, this study assumes that the treatment is definitive rather than empirical: only patients who actually have a bacterial infection are treated, while no considering the uncertainty about the patient diagnosis, which is the main issue we tackle in this research.

The second interpretation of the game-theoretic approach has not yet been considered, and it is the main focus of our study. In this setting, the decision-maker does not completely ignore the effect of antibiotic use on future patients, as in the case of myopic behavior. However, the decision-maker does not completely comply to it as well, as in the case of socially optimal policy. Due to the abovementioned lack of compliance to decision support systems, the Nash equilibrium of this game may be viewed as a realistic representation of physicians' decision-making. Here, the Nash equilibrium is a scenario where none of the physicians benefit from unilaterally changing their behavior, and thus will tend to continue the same course of action. Therefore, it is important to find out under which conditions the equilibrium strategies correspond with the socially optimal policy. A significant difference between the physicians' equilibrium and the optimal policy would lead to the conclusion that full implementation of the social optimum should not rely solely on the goodwill of physicians.

This study analyzes the clinical decision problem of initiating an empirical antibiotic treatment from a game-theoretic perspective and compares the socially optimal policy with the myopic policy and the subgame perfect Nash equilibrium. A subgame perfect equilibrium is the most rational equilibrium of the game, in the sense that each player's strategy is required to be optimal for every possible progression of the game (not only in the path that is actually implemented). In the "Methods" section, the medical problem is described and modeled as a repeated game with imperfect information between two physicians in "Formulation of the problem" subsection, which also includes the description of the myopic policy. The social optimum and its main characteristics are analyzed in "The social optimum" subsection and the subgame perfect equilibrium in "The subgame perfect equilibrium" subsection, where we prove the existence of a symmetric subgame perfect equilibrium and show that the optimal policy and this equilibrium are never the same. Furthermore, we show that the equilibrium strategies always dictate an overuse of antibiotics, compared to the social optimum.

In the "Results" section, we propose a simple yet surprising solution to show that changing the structure of the data available to the physicians, specifically discretizing or coarsening it, can reduce the gap between a rational physician's equilibrium and a socially optimal policy. The formal derivations and proofs of this result are given in "Reducing the problem" subsection of the "Methods" section. This result is independent of any parameter values, apart from the shape of the physician's information distribution (i.e., the distribution of probabilities of a bacterial infection, based on various diagnostics). We show our results under different theoretic distributions as well as an empiric distribution we estimate from data of respiratory infections in children.

Results

The original model: continuous information signal. We describe a repeated (stochastic) game with imperfect information, in which physicians are players, who receive patients over time. Each physician decides whether to treat his current patient with antibiotics. The physicians' decisions rely on a partial information signal about each patient: the probability of a bacterial infection before culture results are available. Each physician estimates this probability, often intuitively, without a formal derivation, from medical symptoms and the results of immediate diagnostic tests[5]. The single-patient decision problem of a physician is described in Fig. 1.
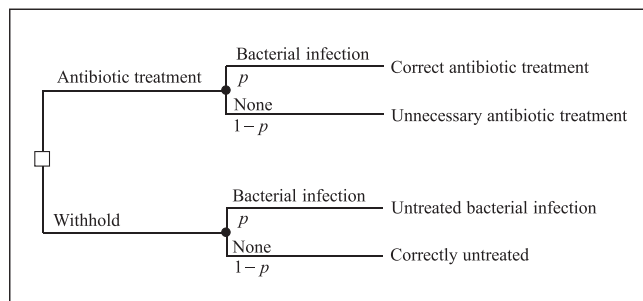
Fig. 1 Single-patient decision tree. The tree of the single-patient decision problem of a physician given information signal p.

In addition, each use of an antibiotic drug depletes its future effectiveness due to antibiotic resistance. Thus, a strategy of the physician in our game describes the minimal probability of infection (threshold signal) that would make him administer antibiotic treatment at any given level of antibiotic effectiveness (E-state—the current sensitivity of the bacterial population to the antibiotic).

In order to calculate the total expected payoff of any possible strategy-combination, the physicians should consider the signal's distribution; i.e., the probability that a future patient would have a certain value of p (probability of a bacterial infection). This information system is characterized by a continuous density function $f(p)$, which induces a posterior p that the patient has a bacterial infection. That is, there is a probability of $F(p) = \int_0^p f(x)dx$ that a random patient has a posterior that is lower than or equal to p.

An extended description of the model appears in "Formulation of the problem" subsection of the "Methods" section.

We then focused on contrasting two main treatment policies: A socially optimal policy, aimed to maximize the sum of all patients' utilities over time; and an equilibrium policy, derived from each physician's attempt to maximize the cumulative utility of all of his own patients over time, given the behavior of the other physician.

Under the social optimum, we show that as long as a sufficient amount of antibiotic effectiveness still exists (the E-state is large enough)—it is worth treating only patients with a very high probability of infection (see Theorem 2 in the section "The social optimum"). When analyzing the equilibrium policy, we prove that a symmetric pure-strategy Markov Perfect Equilibrium (MPE) always exists (Theorem 6 in the section "The existence of a Markov Perfect Equilibrium in pure symmetric strategies"). An MPE is an equilibrium in which each player's strategy is subgame perfect (as explained above), and the decision regarding each patient depends only on the information available to the physician at the present time (p and the current level of antibiotic effectiveness). The decision does not depend, for example, on the past behavior of the other physician. Moreover, under the MPE the physicians always use antibiotics more extensively than the social optimum would dictate (Theorem 7 in the section "Comparing the MPE and the social optimum").

Implementing social optimality. We will try to eliminate, or at least reduce, the gap between the optimal policy and the MPE, by changing the structure of the game. The implementation problem[31] is "the problem of designing a mechanism (game form) such that the equilibrium outcomes satisfy a criterion of social optimality". The mechanism designer cannot impose the socially optimal policy on the players. Instead, he strives to set the rules of the game in such a way that the desired outcome will become an equilibrium. In our case, the possible changes in the rules of the game are constrained, mainly by ethics and norms. For instance,

fining a player for using undesirable strategies is a common and effective means in mechanism design. However, as the payoffs in our game are based on survival rates or other measures of health, it is not possible to punish a physician by reducing his payoffs (i.e., harming his patients). In addition, it is not reasonable to make a clinical decision on the basis of the toss of a coin. Therefore, we will consider solutions that are based on changing the information structure of the game.

Coarsening the information. Our goal now is to manipulate the information structure in a way that would turn the optimal policy into an MPE. Alternatively, in case this is impossible, we would like our manipulation to add a new equilibrium to the game, which is more socially desirable than the existing MPE.

The information system in our original model consists of each physician's private signal, which we define as a continuous density function $f(p)$. This private signal can be, for example, the body temperature of the patient or an output from a decision support system. The continuous signaling system can be replaced with a discrete signal system, which is less informative. The posteriors induced by the new signal are calculated using Bayes' rule. We call this process "coarsening" the information.

If we have a continuous information system $f(p)$ and we want to replace it with a dichotomous discrete system, we need to determine a certain threshold probability T. The new system contains two signals, high (H) and low (L), indicating whether the patient's probability of having a bacterial infection is higher or lower than the given threshold T. Each of these two signals appears with a certain probability and induces a different posterior that the patient has the infection (see section "Coarsening the information: calculations").

Notice that in the case of medical information, it is very simple to implement, especially when decisions are informed by predictive risk scores or probability assessments attained from algorithms—by receiving the results clustered into ranges instead of an exact number. For example, if the signal is the patient's temperature, we can relay information only regarding "high" vs. "low" temperature values, based on our selected threshold.

Similarly, if we want to replace the continuous information system with a discrete system that contains J signals rather than two, we will use a similar coarsening process with $J - 1$ thresholds. However, for the purpose of our discussion, it is sufficient to consider dichotomous signals.

Information and stability. By Theorem 4, we know that if any deviation in any direction (and specifically towards overuse of antibiotics) is worthwhile, then the minimal deviation in this direction is worthwhile as well (the minimal one-stage-deviation principle). However, the opposite claim is not necessarily true. In many cases, it is possible that a physician has an incentive to perform a small deviation towards overuse, while a large deviation in the same direction would result in a loss.

Using a dichotomous signal system with a threshold T is equivalent to limiting the set of available decision rules in each E-state to either "treating only patients with a high signal" or "treating everyone". Consequently, we limited the set of possible deviations from a recommended policy. Specifically, if the recommended symmetric policy is "treating only patients with a high signal", a player may either deviate to "treating everyone" or stick to the recommendation. Thus, it is possible that this recommended policy will not be an equilibrium under the original continuous information system, and yet it will be an equilibrium under the dichotomous system.

As a result, although the optimal symmetric policy is never an MPE in a continuous information system, we may be able to turn it into one by coarsening the information using a very high

threshold. For example, let us assume once again that the signals are the patient's body temperature. Suppose that a body temperature of at least 39 °C induces a very high posterior that the patient has a bacterial infection ($p \rightarrow 1$). Thus, the optimal policy is treating only patients with at least 39 °C, and we want this policy to be an equilibrium. Assume a patient arrives with a body temperature of 38.9 °C. The posterior probability that he has a bacterial infection is only slightly lower than that of a patient whose body temperature is 39 °C, and thus the expected utility from treating this patient with antibiotics is rather high. Therefore, if the physician knows the exact body temperature of this patient, he has a strong incentive to deviate from the optimal policy. However, suppose now that we coarsen the information, and the new signal merely informs the physician whether the patient's body temperature is above or below 39 °C. If the physician only knows that his patient's body temperature is in the range of 36–38.9 °C, the new posterior that this patient has a bacterial infection is much smaller, and therefore the expected utility from treating him is much lower. Consequently, the physician's tendency to follow the optimal policy will be much stronger. A similar scenario may arise if a physician receives an output of a decision support system designed to predict the probability of the need for antibiotic treatment. This output too can be discretized into "low", "medium", and "high" categories instead of a numerical value between 0 and 1 to increase the tendency to follow the optimal policy.

When we use the new information system, the fixed symmetric policy of "treating only patients with a high signal" can be considered a good approximation of the original optimal policy. Therefore, we would like to know whether it can be an MPE.

Our main theorem states that the symmetric strategy combination, in which both players treat only patients with a high signal is an MPE if and only if

$$p_H - p_L \geq [1 - F(T)]p_H + F(T)p_L \qquad (1)$$

For detailed mathematical formulation and proof see Theorem 9 in the section "Implementing optimal policy as a new MPE".

Note that the condition in the theorem involves only the information of the physicians and not any other parameter of the medical problem.

The meaning of Theorem 9 is that treating only patients with a high signal is an MPE if the difference between the higher and the lower posterior of bacterial infection is not smaller than the prior probability of bacterial infection. This prior probability of a bacterial infection, $P(B)$, is a given parameter of the original continuous information system $f(p)$:

$$\forall T, P(B) = [1 - F(T)]p_H + F(T)p_L = \int_0^1 p \cdot f(p)\mathrm{d}p = E(p)$$

which, conveniently, is the expected value of the posterior probability of a patient having a bacterial infection. Therefore, for any given continuous information system, we can define $h(T) = p_H - p_L$, and find out what is the range of possible thresholds, subject to $h(T) \geq E(p)$.

From this theorem, we can also conclude that if we specify the "symmetric strategy combination, in which both players treat only patients with a high signal MPE" as our goal, a coarsening process will fail only in either of the following probabilistic scenarios. First, when the prior probability of bacterial infection (the right side of Eq. (1)) is very high, to begin with. However, in this case, the damage caused by the instability of the optimal policy is limited, since even overuse of the drug would yield a relatively high social utility. Second, if the original signal distribution has low dispersion resulting in a low difference between p_H and p_L (the left side of Eq. (1)).

Signal distributions. To demonstrate the relations between the original information structure and the possible values of the threshold T, we present several possible scenarios of the signal distributions of a bacterial infection $f(p)$ represented by a beta distribution with different parameter combinations (Fig. 2a–d). The choice of a beta distribution was driven by its flexibility, its support laying in [0,1], and its role as a conjugate prior of the Bernoulli distribution, lending itself naturally to model distributions of such probability values. Examples of other distributions are given in the Supplementary Information (Supplementary Note 12). First, we present a scenario wherein often information regarding the infection is distinctly associated with a bacterial or non-bacterial disease (Fig. 2a). Such a scenario may arise when diagnosing endocarditis, where clinical information can be highly indicative of bacterial infection, rule it out almost certainly, or have some uncertainty[32]. In this case, we can set T at any value larger than 0 (Fig. 2e). In a scenario where patients are approximately as likely to arrive with different symptoms that define
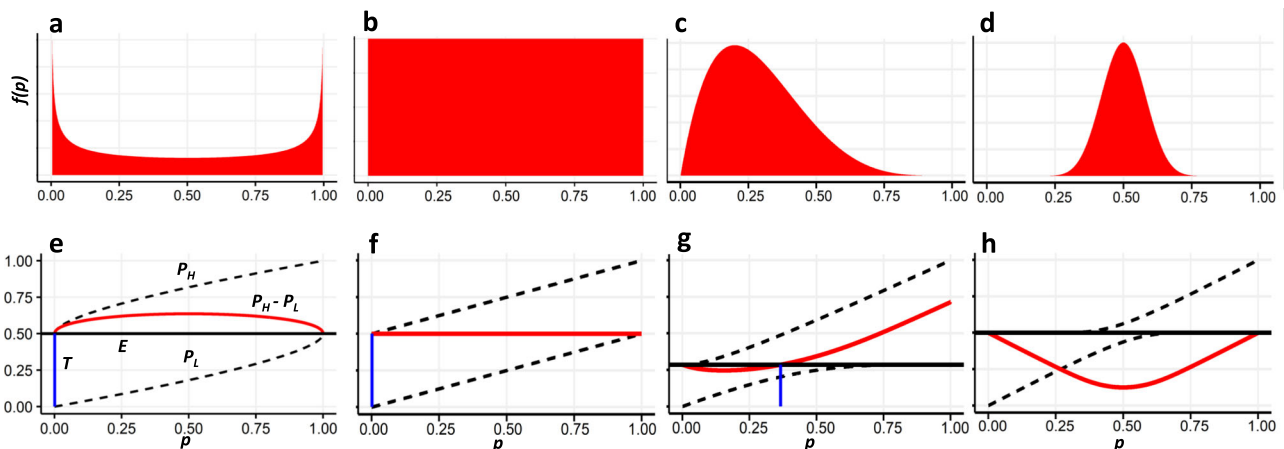

Fig. 2 Theoretical posterior distributions of bacterial infection signals, and corresponding threshold values possible by coarsening the data.
a–d Probability distribution functions $f(p)$ representing distinct scenarios of a bacterial infection signal. **e–h** Corresponding (top-bottom) values of E (expected posterior value, horizontal black lines), p_H and p_L (posterior probability of bacterial infection given a high and low signal, upper and lower dashed curves, respectively), their difference $p_H - p_L$ (red curves) and the minimal threshold T (blue vertical lines) achievable by coarsening the data. All distributions are beta with parameters $\alpha \in \left(\frac{1}{2}, 1, 1, 20\right)$, $\beta \in \left(\frac{1}{2}, 1, 5, 20\right)$, corresponding to **a–d**.

various degrees of certainty in the diagnosis of a bacterial infection, we can represent $f(p)$ by a uniform distribution (Fig. 2b). Similarly, we can model information with uniform distribution on only a subset of the range of probability values, e.g., uniform on (0,0.5) and zero otherwise. A scenario pertaining to these kinds of distributions can be observed, for example, in pharyngitis diagnosis, for which a set of criteria have a range of probabilities attributed to them being indicative of a bacterial infection[33]. Then, once again, any threshold T can be turned into an MPE by coarsening the information (Fig. 2f). If most of the infections are non-bacterial, and $f(p)$ is skewed to the right, we have some freedom in setting the value of T and can potentially reduce antibiotic misuse by coarsening the signal given to physicians (Fig. 2c). This is similar to the empiric data of upper respiratory infections in young children, demonstrated in the next section. On the other hand, if in most cases a bacterial infection is approximately as likely as a non-bacterial one, i.e., $f(p)$ is centered around 0.5, then there is not enough information in the system and no equilibrium T values can be set by coarsening the data (Fig. 2d).

Empirical example. We further investigated a real-world example. We obtained medical record data of 1202 children aged <2 years who were hospitalized for a respiratory tract infection. We trained a machine learning model to classify bacterial and viral infections using the patients' medical records, and used the model's output as an estimate of the posterior probability of a bacterial infection—$\hat{f}(p)$ (Fig. 3; see "Methods" section for details). For the obtained posterior distribution, coarsening the data can produce any threshold $T \geq 0.546$, which can thus allow varying degrees of the strictness of antibiotic regulation upon a respiratory tract infection in children. That is, if the output of this toy model, or an improved machine learning algorithm, is provided to physicians to assist rapid classification of respiratory tract infections, a simplified output of "high" and "low" probability should help compliance and might mitigate antibiotic overuse.

Discussion

This study applies the conceptual framework of game theory to analyze the conflict between individual and societal interests that is inherent to the decision problem of initiating an empirical antibiotic treatment. The problem is clear when the physician focuses on the well-being of a single present patient: in this case, initiating antibiotic treatment is a dominant strategy. However, we have shown that the problem remains even if the physician does consider long-term implications, and cares about the well-being of his future patients.

We have shown that in this setting, although a symmetric subgame perfect equilibrium in pure Markov strategies always exists, it is not socially optimal, and the socially optimal policy is not an MPE. The analysis in the section "The subgame perfect equilibrium" of the "Methods" chapter, provides evidence that, although administering antibiotics to everyone is no longer a dominant strategy, the rational strategic behavior would lead to an overuse of antibiotics.

Our model reveals that in a setting that involves an interaction between several decision-makers (physicians), supplying each of them with the maximal available information turns the socially optimal treatment policy into an instable recommendation (i.e., knowing the maximal information generates an incentive for each physician to deviate from the recommended policy). This is an unintuitive result since in single-agent decision theory information cannot have a negative value, as it is either beneficial or ignored by the decision-maker. We have further suggested a way to mitigate the problem by coarsening the information that is
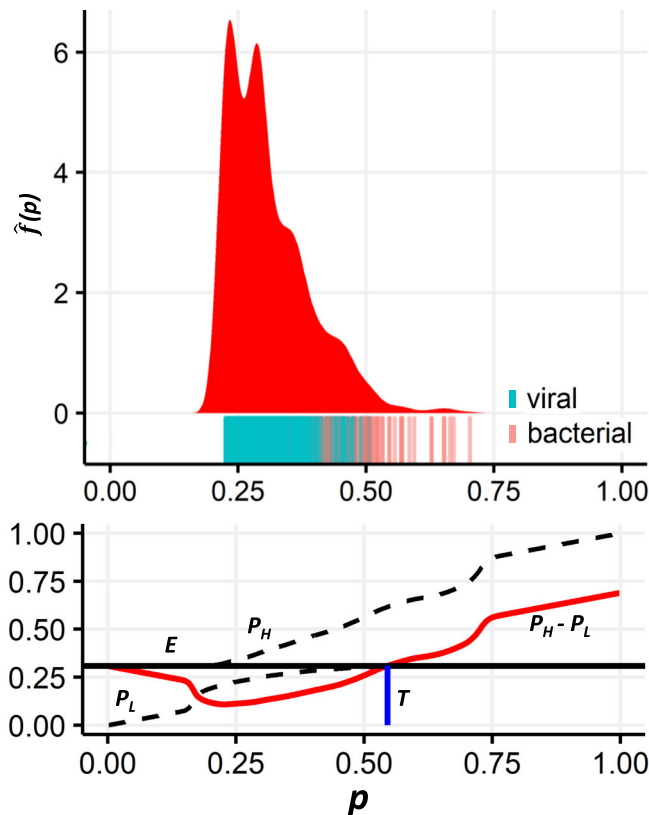

Fig. 3 Posterior distributions of bacterial infection signals for children with respiratory tract infections, and the corresponding threshold values possible by coarsening the data. Top: The estimated probability distribution function $f(p)$ associated with respiratory tract infections in children, caused either by a bacterial or viral (RSV) agent (see "Methods" section). Under the distribution, a rug-plot indicates the actual origin of infection in the data. Bottom: The corresponding values of E (expected posterior value, horizontal black line), p_H and p_L (posterior probability of bacterial infection given a high and low signal, upper and lower dashed lines, respectively), their difference $p_H - p_L$ (red curve) and the minimal stable threshold T (blue vertical line) achievable by coarsening the data.

available to the physicians regarding the probability of a bacterial infection. Hence, by coarsening the available information and bringing the physician's decisions closer to the social optimum, a reduction in antibiotic misuse is expected, which should consequently reduce antibiotic resistance frequencies[1,11,13,34,35].

Importantly, the conditions dictating when coarsening the information is socially beneficial to depend solely on the distribution of the signal indicating whether the infection is of bacterial origin. This is a general result, independent of other parameters such as the utility of antibiotic treatment and antibiotic resistance frequencies.

Coarsening the information provided to a physician is already carried out in both this and other medical decision problems. For example, the Centor criteria help identify whether a sore throat is caused by a bacterial agent[33], and the diagnosis is based on summing points allotted to symptoms and discretized age brackets. Similar discretizations of risk information are made for stroke[36], endocarditis[32], cardiovascular-related death[37], pancreatic fistulas[38], and many other medical applications. Our work shows for the first time that this discretization can guide rational physicians to more societally beneficial decisions.

Furthermore, low compliance to the recommendations provided by clinical decision support systems is often a limiting factor in their success[39,40]. In addition to the benefits of coarsening information

presented in our results, a potential benefit could result from the fact that providing detailed probabilistic information to physicians does not necessarily improve their decisions regarding antibiotic treatment. A study examining whether improving physicians' judgments of the probability of streptococcal pharyngitis for patients with sore throats would affect their use of antibiotics, concluded that teaching physicians to make better judgments of disease probability may not alter their treatment decisions[5]. The authors hypothesized that even when the exact probabilistic information was available to physicians, they did not condition their decisions on it; but rather on ordinal judgments of disease probability or symptom severity, such as "not sick", "mildly sick", "moderately sick", or "severely sick". If this is the case, coarsened information can define these categories for physicians in a way that would benefit society. In other words, the goal of the coarsening is not to reduce the information available to the physicians but to simplify the decision-making process and thus make it more effective.

Although the analysis performed here is relatively general and relies on few assumptions, it naturally includes several limitations. Principally, the coarsening process may also add new undesired equilibria. In the worst case, it may even turn the myopic policy of treating everyone at every E-state into an equilibrium. However, a reasonable assumption is that a policy can be recommended to the physicians (either by some kind of authority or by a decision support system), and if the players have no incentive to deviate from this recommendation—it will indeed be accepted and implemented. Therefore, if the coarsening process indeed succeeds, the social interest in adding a good MPE outweighs the risk of adding a bad MPE.

Notably, we did not use a mechanistic model for the spread of antibiotic resistance, but rather approximated it by depletion of antibiotic effectiveness (E-state), as explained in the "Methods" section. We also assumed, for mathematical convenience, that the game ends when antibiotic effectiveness is completely depleted (E-state 0). In practice, the administration of antibiotics with very high resistance frequencies is rare, but this has a negligible effect on our modeling framework. Low levels of drug effectiveness yield very low utility to the players, and therefore stopping the game at this point has no significant effect on their decisions until this point. Moreover, while we set the conditions for a successful information coarsening process (Theorem 9), we demanded that the new threshold T would be an equilibrium at any resistance level. If we, alternatively, specify a less ambitious MPE as our goal, e.g., "treating only patients with a high signal from E-state k' onwards" (i.e., restricting the antibiotic treatment to patients with likely bacterial infections only when resistance levels are low), such a target MPE would only extend the range of possible T values.

Further limitations of our analysis are the focus on a decision including only the binary option of treatment/no treatment, and focusing on a two-player game between physicians. Regarding the former, a more general scenario should include different types of antibiotics, with different spectra of coverage and effect on resistance frequencies. In order to simplify the mathematical presentation, we have chosen to focus on the single antibiotic case, which contains the main aspects of the decision problem. Future research should also cover a wider framework, containing several drugs, e.g., representing the prescription of broad or narrow-spectrum antibiotics. However, this is a more complex case where the effect of one antibiotic on the other's future effectiveness, due to cross-resistance[41–43] between the different antibiotics, is not trivial and should be accounted for. With respect to the latter limitation, our results are relevant not only for single physicians. A player in a game-theoretical model is often a group of people with common goals and utilities and centralized decision-making, such

as a football team, a firm, or even a country. In our case, a player may refer to a hospital or a country, in which many physicians apply the same clinical guidelines and policies. Consequentially our results can translate into the conflicts between the current and future patients of the specific hospital or country, rather than a certain physician. Finally, we note that our results are of the highest relevance where antibiotic use is mostly conditioned on a physician's prescription. We acknowledge that this might not be the case in some low- and middle-income countries, where non-prescription use can even account for the vast majority of total antibiotic use[44].

To conclude, the problem of antibiotic misuse and resistance can be viewed as an example of the tragedy of the commons[45]. This concept has been previously used to describe the allocation of limited medical resources in scenarios of a conflict between individual and societal interests[46]. Previous studies called for collective solutions[47], which may take the form of placing explicit resource constraints on resources available to physicians, or clinical practice guidelines that recognize cost-effective care as acceptable. The current study suggests another type of solution: supplying the physician with the level of information that would enable him to be committed to his patients without hurting society. This solution is expected to become increasingly relevant as data collection and computational power continually enhance the involvement of decision support systems in medicine[48] and specifically with respect to antibiotic resistance[49–51], and hence warrants further empirical investigation.

Thus, by changing his available information, we allow the physician to act rationally and yet guide him towards the socially optimal decision and relieve him of the prisoner's dilemma underlying the decision of antibiotic prescription.

Methods

Formulation of the problem. The medical problem considered in this research involves a patient who has clinical symptoms that may indicate the existence of a bacterial infection. To simplify the analysis, we assume that only one type of antibiotic treatment is available, and that antibiotic treatment increases the chances of recovery of patients having bacterial infections while not affecting the recovery of patients without bacterial infections.

The effectiveness of the antibiotic drug, denoted e, is the probability that the infecting bacteria are susceptible to it, whereas $1 - e$ will be the probability of resistance.

The effectiveness of the antibiotic drug is not constant. Each use of the antibiotic drug (whether necessary or not) induces selection of resistant pathogens, and therefore reduces its future effectiveness (i.e., diminishes e). In addition, a patient who does not have the bacterial infection might acquire an opportunistic infection as a consequence of unnecessary antibiotic treatment. The chance of that happening is denoted c.

The physician, an agent acting on behalf of the patient, must decide whether an empirical antibiotic treatment should be administered. His decision relies on partial information about the patient: the probability of a bacterial infection in a given patient before culture results are available. This information is derived from symptoms, immediate diagnostic tests or decision support systems. The information signal results in a random variable parameter p, which is the posterior probability that the patient has the bacterial infection, where $0 \leq p \leq 1$. The signal is a continuous random variable with a density function $f(p)$ (see section "Treatment policies in the dynamic model", under "Information"). We assume the physician has complete information on all the parameters mentioned above, regarding the patient.

Treatment policy in the static model. Most of the medical decision-making focuses on the best interest of a single patient currently under treatment. Under this point of view, the question of whether to administer an antibiotic treatment, given a certain signal, becomes a static decision problem. The static decision analysis takes the current level of effectiveness as a parameter and ignores the current decision's contribution to the emergence of resistance and its effects on other future patients.

Following Pauker and Kassirer[52,53], the static problem can be represented as a decision tree, and analyzed using the threshold approach derived directly from expected utility theory. In our case, the decision tree is presented in Fig. 1. The calculation of the expected utility of each action is simply performed by weighting the probability and the utility of each possible outcome presented in this tree. If we denote bacterial infection by B and no bacterial infection by N, the expected

utilities of administer\withhold antibiotic treatment are:

$$E(A) = p \cdot u(B, \text{Treated}) + (1-p) \cdot u(N, \text{Treated})$$
$$E(W) = p \cdot u(B, \text{Not Treated}) + (1-p)$$
$$\cdot u(N, \text{Not Treated})$$

where p is the probability of having a bacterial infection and u is the utility function. The aim of the static analysis is to calculate the treatment threshold probability, T. The optimal treatment policy of the physician is administering treatment if the probability of the patient having a bacterial infection (p) exceeds the threshold. By setting $E(A) = E(W)$, the treatment threshold probability is calculated in the following way:

$$T = \frac{u(N, \text{Not Treated}) - u(N, \text{Treated})}{(u(B, \text{Treated}) - u(B, \text{Not Treated})) + (u(N, \text{Not Treated}) - u(N, \text{Treated}))} \tag{2}$$

In the original analysis of Pauker and Kassirer, there are four possible outcomes, and each one of them is assigned a utility. However, as we have mentioned, in our case we assume that a patient who does not have the bacterial infection has the same chance of recovery with or without antibiotic treatment. Furthermore, we assume that since the case of bacterial infection involves a substantial risk, treating bacterial infection is always preferred to not treating it, and not having a bacterial infection is preferred to having it, regardless of treatment. Thus, we can define three levels of utility: r_1, the utility of "Bacterial Infection - Not Treated" (bacterial infection, without appropriate empirical antibiotic treatment); r_2, the utility of "Bacterial Infection - Treated" (bacterial infection, with appropriate empirical antibiotic treatment); and r_3, the utility of "No Bacterial Infection" (no bacterial infection, either with or without antibiotic treatment). Where $r_3 > r_2 > r_1$ (note that the assumption regarding r_3 is presented at this point only for the sake of compliance with the approach of Pauker and Kassirer, but it will not be needed later on in our complete model, from Eq. (5) onward).

For example, if we assume that the most important aspect of the possible outcomes is the survival of the patient, the utilities were assigned as follows: There are two possible states—alive and dead. Since we are dealing with Von Neumann–Morgenstern utilities[54] we can arbitrarily set $u(\text{alive}) = 1$ and $u(\text{dead}) = 0$. Consequently, the utility of each outcome is simply the conditional probability of survival. Importantly, the utilities in the model can include other aspects not mentioned above, such as drug side effects, quality of life, financial costs, or other indirect costs[6,55]. The only assumption we make is that the order of preferences, based on the utility values, is preserved ($r_3 > r_2 > r_1$). This assumption translates into the statement: having an untreated bacterial infection is worse than having a bacterial infection that is treated with an effective antibiotic, and this is, in turn, is worse than not having a bacterial infection at all. This assumption holds as long as

having a bacterial infection is costly (health-wise or financially) to the patient or the healthcare system.

If we replace the final outcomes in the original tree with these utilities, denote A for administer treatment and W for withhold treatment, and add the posterior probability of disease induced by the patient's symptoms (p), we get to the decision tree shown in Fig. 4a.

However, two additional components of our model need to be included: the current level of drug effectiveness (e), and the risk that a patient who does not have the bacterial infection will develop one because of antibiotic treatment (c). These two components have the form of additional chance nodes in our decision tree. Importantly, these do not add new type of outcomes, since they are simply lotteries between the worst outcome and a more preferable one. The decision tree of the single-patient one-period problem given a signal that induces a posterior p is shown in Fig. 4b. The definitions of all relevant variables are summarized in Table 1.

If we now want to calculate the threshold treatment probability in our model, Eq. (2) has to be adjusted in accordance with the new decision tree. The threshold T is defined as the probability (of bacterial infection) that generates indifference between administering and withholding antibiotic treatment. The optimal treatment policy of the physician is administering antibiotic treatment if the conditional probability induced by the patient's signal exceeds the threshold: $T \leq p$. The net risk of treatment is the loss of utility (or survival rates) caused by developing a bacterial infection as a consequence of the antibiotic treatment, multiplied by the chance of that happening. The net benefit of treatment is the net gain of utility (or survival rates) due to effective antibiotic treatment, multiplied by the chance that the treatment is indeed effective. The threshold treatment probability equation becomes:

$$T = \frac{c(r_3 - r_1)}{e(r_2 - r_1) + c(r_3 - r_1)} \tag{3}$$

Note that unlike Eq. (2), in the adjusted threshold Eq. (3) there are only three possible final outcomes (instead of four) and two extra multiplications: of the net benefit by the effectiveness coefficient e and of the net risk by the risk coefficient c.

Treatment policies in the dynamic model. Treatment policies in dynamic models involve the use of an antibiotic drug over time. In our dynamic setting, drug susceptibility (effectiveness) is regarded as a nonrenewable resource. This is a decent approximation of the state of antibiotic susceptibility—reversal of resistance is relatively slow and works on time scales bigger than resistance accumulation[56,57]. Therefore, any policy which aims to maximize the utility of different patients over time (either all patients or a selected group of them) has to take into consideration the future negative consequences of antibiotic use. For the purpose of calculating social utility, we will assume that it equals the sum of the individual utilities.

The dynamic model compares three types of treatment policies: a myopic policy, aimed to maximize the utility (or chance of survival) of the current patient
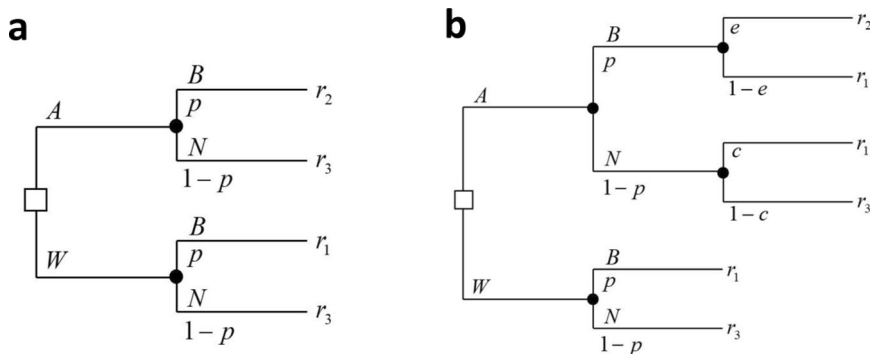
Fig. 4 The decision tree of the single-patient static problem given signal p. a The basic decision tree including administering or witholding treatment (A and W, respectively), with the probabilities of bacterial and non-bacterial infections (B and N, respectively), and their associated utilities of the outcomes (r_1, r_2, r_3). **b** The tree extended to include antibiotic effectiveness (e) and the risk of unnecessary treatment (c). All variable definitions are given in Table 1.

Table 1 Definitions of the variables in the single-patient static problem.

A	Administer antibiotic treatment
W	Withhold antibiotic treatment
p	Probability of bacterial infection given the symptom (signal)
e	Current resistance level of the antibiotic drug
c	Chance of developing bacterial infection by resistant pathogen after using the antibiotic drug
r_1	Utility of bacterial infection, without appropriate empirical antibiotic treatment
r_2	Utility of bacterial infection, with appropriate empirical antibiotic treatment
r_3	Utility of no bacterial infection, either with or without antibiotic treatment

at each point of time; a socially optimal policy, aimed to maximize the sum of all patients' utilities over time; an equilibrium policy, derived from each physician's attempt to maximize the cumulative utility of all of his own patients over time, given the behavior of the other physician.

Time in our dynamic model is discrete, each physician receives patients and makes decisions at discrete time points. The following definitions specify the characteristics of the game model.

Players. The basic setting is a game of two players (physicians) denoted 1 and 2, each of whom treats one patient in each period. A brief discussion of the extension to n players is given in section "Recursive calculation of payoffs" and a more general interpretation of the two players model appears in the discussion.

Patient's health condition. A patient either has a bacterial infection, B, or does not have it N. The patient's condition is not known to his/her physician. We denote by h_i the health condition of player i's patient, $h_i \in \{B, N\}$, $i = 1, 2$. The prior probability of a patient having the bacterial infection in each period is $P(B)$. These probabilities are independent within each period (between the two patients) and between periods (over time).

Actions. Each player chooses one of two possible actions: to administer antibiotic treatment, A, or withhold antibiotic treatment, W.

Effectiveness depletion. As mentioned before, each use of the drug decreases the average antibiotic susceptibility to the drug, or its effectiveness, in the next period by a depletion effect α. Such depletion is in accordance with both biological knowledge of antibiotic resistance acquisition in bacterial populations as well as observations of resistance frequencies in countries with varying antibiotic consumption[11,34,35]. The game ends when antibiotic effectiveness is completely depleted, for mathematical convenience. This assumption can be easily attenuated (see "Discussion" section). Given an initial effectiveness level e_0 and the depletion effect α, one can calculate the total number of possible game effectiveness-states, M. Note that since the effect of use has a delay of one period, when the effectiveness is α two patients can be treated simultaneously before the effectiveness is completely depleted. Therefore, the number of patients that will be treated is either M or $M + 1$. Empirically, M is very large (an antibiotic drug is used by millions of patients before it is considered not effective), and for the purpose of mathematical convenience, we will assume that it is as large as we want. For simplicity, we will also assume that $e_0 = M\alpha$, where M is a positive integer. The current level of effectiveness e_t in any given period t is represented by the current effectiveness-state (E-state), $k_t = \frac{e_t}{\alpha}$. For simplicity, and since the E-state does not depend on the period number, we will omit the index t and denote the E-state by k.

States and dynamics. The definition of state in our model has two aspects. As explained, the current E-state is defined by the current number of doses left, $k = \frac{e_t}{\alpha}$. In each E-state, the players face a stage-game with imperfect information. The information structure is defined in the next paragraph. In each stage-game, there are four health-states (H-states), which are the combinations of the health condition of each of the current patients. The set of H-states is $\Omega = H_1 \times H_2$, where $H_i = \{B, N\}$ represents the health condition of player i's patient. That is, for any fixed k the four possible H-states are: $\{(B, B), (B, N), (N, BB), (N, N)\}$. The system dynamically moves between states. The system dynamics has two components: deterministic and stochastic. The deterministic component is the transition between E-states, and it is determined by the player's actions and the current E-state number k, through the effectiveness depletion. The stochastic component is related to the condition of the patients in the next period. It is stationary and depends on the distribution of patients, i.e., on $P(B)$. Thus, with $k > 1$ doses left, the deterministic dynamics is moving from E-state k to k, $k - 1$ or $k - 2$ according to the actions described in Table 2.

The stochastic dynamics are:

$$P(\omega) = \begin{cases} (P(B)^2) & \text{for } \omega = (B, B) \\ (1 - P(B))^2 & \text{for } \omega = (N, N) \\ P(B)(1 - P(B)) & \text{for } \omega = (N, B) \text{ or } \omega = (B, N) \end{cases}$$

Table 2 Transition matrix of the deterministic game.

	Player 2's action	
	A	**W**
Player 1's action		
A	$k - 2$	$k - 1$
W	$k - 1$	k

A administer treatment, W withhold treatment, k doses left (E-state).

Information. We assume that each physician knows the E-state k. That is, physicians have information regarding the average resistance frequencies of bacterial infections in their cohorts. However, since the health condition of the patients is not known, the current H-state is not known by any physician. Each physician observes a signal regarding his patient. The information signal results in a random variable parameter p, which is the posterior probability that the patient has the bacterial infection, where $0 \le p \le 1$. The parameter p is a continuous random variable with a density function $f(p)$. We assume that f is integrable, and thus the probability that a patient's posterior is p' or less is $F(p') = \int_0^{p'} f(x) dx$. For simplicity, we will also assume that $f(p) > 0, \forall p$. This assumption can be easily omitted, by adjusting the strategy space, as explained later under "Decision rules and strategies" section. The patients' signals within and between each period are independent. The signals are private information; each physician knows only his own patient's signal. The distribution of the signals, $f(p)$, is common knowledge. Note that the prior probability that a patient has the bacterial infection is: $P(B) = \int_0^1 p \cdot f(p) dp$, which is the expected value of the posterior distribution of a bacterial infection, later referred to as $E(P)$.

The immediate payoffs. In each period these are defined as the net expected utility of antibiotic use, i.e., the difference between the utility of treatment and the utility of no treatment (otherwise, infinite utility can be accumulated by not treating any patient forever, due to the chance of a spontaneous recovery). Positive net utility (or gain in survival chances) can be achieved only if the patient had a bacterial infection and the treatment was effective. The probability of that happening, when signal p_i is observed, is $p_i \cdot e = p_i \cdot k\alpha$. Negative utility (or loss of survival chances) is the consequence of developing infection as a result of the antibiotic treatment, when the patient did not initially have the infection. The probability of that event is $c \cdot (1 - p_i)$. Following are the immediate expected payoffs of physician i in E-state k when the signal of his patient is p_i:

If physician i plays A:

$$u_{p_i}(k) = p_i k\alpha(r_2 - r_1) - (1 - p_i)c(r_3 - r_1) \qquad (4)$$

If physician i plays W: 0

Notice that the immediate payoff of each physician depends only on his own action in the current period. The mutual effect is indirect and delayed, through the depletion of e.

Decision rules and strategies. A strategy of a physician is a mapping from histories (of E-states, H-states, and actions) to actions. However, we are interested in subgame perfect equilibrium in stationary Markovian strategies (also known as Markov perfect equilibrium—MPE). Markovian means dependence only on payoff-relevant variables, which in our model are the states (the E-states and the partial information about the H-states). Stationary means that these strategies will not depend on t (since time is not payoff-relevant in our model). In order to check that a combination of stationary Markovian strategies is an equilibrium, we need only to check that each player has no incentive to deviate to another Markovian strategy. The reason is that given any fixed stationary Markov strategy played by the other physician, the decision problem faced by physician i is equivalent to a Markov decision process (MDP)[58]. Thus, the best response exists in Markov strategies, and it can be found using a maximization process of dynamic programming. Therefore, we will denote strategies as Markovian, even though we do not actually limit a physician from deviating into a non-Markov strategy. A stationary Markov strategy is compounded of decision rules. A decision rule of a physician determines what to do in the current E-state, given the signal that he currently observes (the physician's information about the current H-state), and not on the history. We will limit our discussion to threshold decision rules $d_i^k \in [0, 1)$, where physician i will choose A if $p_i \ge d_i^k$ and will choose W if $p_i < d_i^k$. The decision rule $d_i^k = 1$ is not allowed, because it means not treating at all, since the probability that a patient has a posterior of $p_i = 1$ or more is $1 - F(1) = \int_1^1 f(x) dx = 0$. For the same reason, if we want to allow $f(p) = 0$ for $p_{max} < p \le 1$ then we must limit $d_i^k \in [0, p_{max}]$. A pure (stationary Markov) strategy of physician i is, therefore, a vector containing the physician's M decision rules for all the possible E-states $s_i = (d_i^1, d_i^2, \ldots, d_i^k, \ldots, d_i^M), i = 1, 2$. Note that the decision rules will be implemented in reverse order (d_i^1 is the decision rule of player i with one dose left to give, $e = \alpha$). In addition, note that not all the rules will necessarily be applied in the realization of the game, because whenever both physicians choose A simultaneously the game skips E-state $k - 1$ and moves straight to $k - 2$. We denote by $s_i^k = (d_i^1, d_i^2, \ldots, d_i^k)$ (the projection of s_i on its first k coordinates) a strategy in the subgame that starts on E-state k (when $\frac{e_t}{\alpha} = k$) and continues onwards. We note that threshold decision rules are not only the intuitive, it is also easy to prove that they are the most efficient. The proof requires additional definitions that will be specified later in this section and appear in the Supplementary Information (Supplementary Note 3).

Finding the myopic policy is immediate. The emergence of resistance changes the net benefit of action A (giving antibiotics) over time, but the future change has no influence on the decision in the current period. Therefore, the threshold probability is calculated in each period using (3) with the current level of effectiveness, e_t.

However, in order to find the equilibrium policies and the socially optimal policy, we need to analyze the repeated game characterized by the model. The socially optimal policy is derived assuming that one physician (a social planner) treats all the patients. Under this assumption, the game becomes a Markov decision process.

The information structure of the social planner problem matches the information structure of the game between the two physicians. Namely, the social planner is only allowed to condition his treatment policy on each patient's signal, and not on the combination of signals.
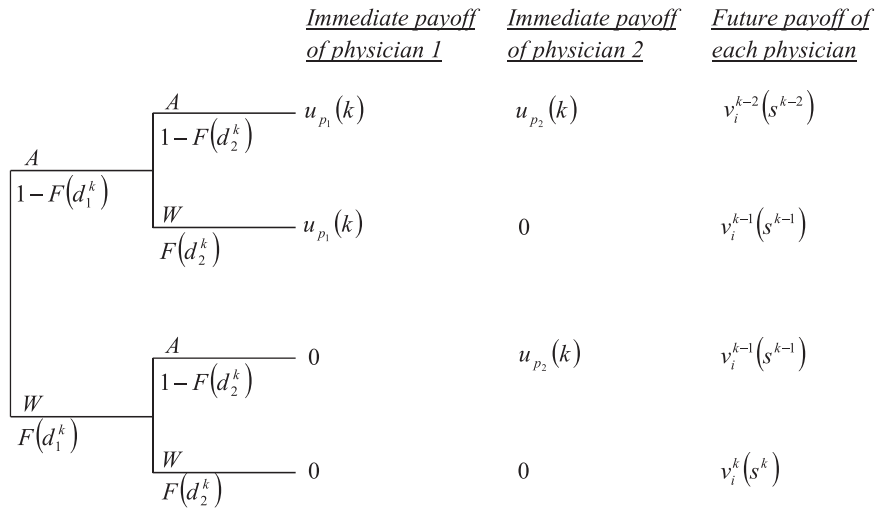
	Immediate payoff of physician 1	Immediate payoff of physician 2	Future payoff of each physician

Fig. 5 Expected payoffs of the players on E-state k. All variables are defined in Table 3.

This setting enables us to compare the two cases since it preserves the information limitations of the game, but it is also reasonable—in reality, the number of both patients and physicians is much higher, and it is very unlikely to allow an action that depends on the simultaneous information of all. Following this line of reasoning, we will also require that the socially optimal policy will be symmetric, i.e., the same decision rule (but importantly, not necessarily the same decision) is applied to both patients in every period.

The analysis of the game requires a more detailed definition of the payoffs and the concept of deviation. These definitions and the derivation of both optimal and equilibrium policies are described in the next section. For the sake of simplicity, and since c is very small, we will assume from this point onwards that $c = 0$. Under this assumption, the net risk of treatment equals 0, and using antibiotics is a dominant alternative from the single-patient perspective. As a result, the myopic policy is necessary to treat every patient, regardless of his signal. The immediate expected payoff for a player who chooses action A in E-state k under this assumption is:

$$u_{p_i}(k) = p_i k \alpha (r_2 - r_1) \tag{5}$$

Recursive calculation of payoffs. Payoffs can be calculated for any strategy profile $s = (s_1, s_2)$. Let $v_i^k(s^k)$ be the expected cumulative payoff of physician i from E-state k onwards, for a strategy profile s. The expected payoff is calculated with respect to the probability distribution of the signals. Note that the payoff depends on the strategies of the players only from this point onwards (i.e., their decisions for E-states $1, \ldots, k$), and on the distribution of the future signals.

The value of $v_i^k(s^k)$ can be calculated recursively, starting from $k = 1$. The action taken by each physician in each E-state provides him with a certain immediate payoff, and the combination of actions taken by both physicians determines the future E-state, and thus their future payoff. Given the current E-state decision rule (threshold) of each physician, d_i^k, the probability that physician i will choose to administer treatment equals the probability that his patient's signal will exceed the threshold, i.e., $\int_{d_i^k}^1 f(p)\mathrm{d}p = 1 - F(d_i^k)$. The different possibilities of expected payoffs on E-state k are represented in Fig. 5, and the definitions of all relevant variables appear in Table 3. Note that the figure does not represent the game tree, but merely the four possible action combinations and their consequences.

To facilitate the recursive calculation of payoffs, we will sometimes write $v_i^k(s^k)$ as $v_i^k(s^{k-1}, d_i^k, d_{-i}^k)$. The payoffs can be calculated using the following recursive equation:

$$v_i^k(s^{k-1}, d_i^k, d_{-i}^k) = \int_{d_i^k}^1 f(p)u_{p_i}(k)\mathrm{d}p + F(d_i^k)F(d_{-i}^k)v_i^k(s^k)$$
$$+ \left[F(d_i^k)(1 - F(d_{-i}^k)) + (1 - F(d_i^k))F(d_{-i}^k)\right]v_i^{k-1}(s^{k-1})$$
$$+ (1 - F(d_i^k))(1 - F(d_{-i}^k))v_i^{k-2}(s^{k-2})$$

or, noting that $v_i^k(s^k)$ is on both sides of the equation,

$$v_i^k(s^{k-1}, d_i^k, d_{-i}^k) = \left(k\alpha(r_2 - r_1)\int_{d_i^k}^1 p \cdot f(p)\mathrm{d}p + \left[F(d_i^k)(1 - F(d_{-i}^k))\right.\right.$$
$$\left.+ (1 - F(d_i^k))F(d_{-i}^k)\right]v_i^{k-1}(s^{k-1})$$
$$\left.+ (1 - F(d_i^k))(1 - F(d_{-i}^k))v_i^{k-2}(s^{k-2})\right) \cdot \frac{1}{1 - F(d_i^k)F(d_{-i}^k)}$$
$$\tag{6}$$

Table 3 Definitions of the payoff variables in the dynamic game.

A	Administer antibiotic treatment
W	Withhold antibiotic treatment
k	Current E-state (number of antibiotic doses left)
d_i^k	The threshold decision rule of physician i
$1 - F(d_i^k)$	The probability that physician i administers treatment
$F(d_i^k)$	The probability that physician i withholds treatment
$u_{p_i}(k)$	The immediate payoff that physician i receives from treating the current patient
$v_i^{k-j}(s^{k-j})$	The future payoff of each physician, when exactly j of the current patients are given antibiotics

The numerator in (6) is the expected immediate payoff of player i given his own current E-state decision rule plus the expected future payoff given the strategy profile. The expected future payoff is the probability that either one or two physicians give antibiotic treatment, multiplied by the expected payoff with either one or two doses left, respectively. The denominator is the probability that at least one physician gives antibiotic treatment. For the sake of simplicity, we will sometimes use the following notation:

$k \cdot U = k\alpha(r_2 - r_1)$—the maximal immediate payoff of each physician (net benefit of treatment) at E-state k.

$E_i(d_i^k) = \int_{d_i^k}^1 p \cdot f(p)\mathrm{d}p$—the expected posterior of a patient that is treated by physician i.

$A_j(d_i^k, d_{-i}^k)$—the probability that exactly j physicians treat with antibiotics. Specifically:

$$A_0(d_i^k, d_{-i}^k) = F(d_i^k)F(d_{-i}^k)$$
$$A_1(d_i^k, d_{-i}^k) = F(d_i^k)(1 - F(d_{-i}^k)) + (1 - F(d_i^k))F(d_{-i}^k)$$
$$A_2(d_i^k, d_{-i}^k) = (1 - F(d_i^k))(1 - F(d_{-i}^k))$$

Using this notation, (6) can be written as

$$v_i^k(s^{k-1}, d_i^k, d_{-i}^k) = \frac{k \cdot U \cdot E_i(d_i^k) + A_1(d_i^k, d_{-i}^k)v_i^{k-1}(s^{k-1}) + A_2(d_i^k, d_{-i}^k)v_i^{k-2}(s^{k-2})}{1 - A_0(d_i^k, d_{-i}^k)} \tag{7}$$

Remark. (6) can be easily extended to n physicians:

$$v_i^k(s^k) = \frac{k\alpha(r_2 - r_1)\int_{d_i^k}^1 p \cdot f(p)\mathrm{d}p + \sum_{j=1}^n P(\text{exactly } j \text{ physicians treat})v_i^{k-j}(s^{k-j})}{1 - \prod_{j=1}^n F(d_j^k)}$$

Note that since $F(d_i^k)$ is continuous and $0 \leq d_i^k < 1$, the payoff function $v_i^k(s^{k-1}, d_i^k, d_{-i}^k)$ is also continuous. Due to the symmetry between the players and the symmetry of the social optimal policy, we will concentrate on searching for symmetric Nash equilibria. Therefore, we will assume from this point onwards that the strategy profile is symmetric, i.e., $d_i^k = d_{-i}^k = d^k$ for all $k = 1, \ldots, M$. Under the

symmetry assumption, (6) can be replaced by:

$$v_i^k\left(s^{k-1}, d^k, d^k\right) = \left(k\alpha(r_2 - r_1)\int_{d^k}^1 p \cdot f(p)\mathrm{d}p + 2F\left(d^k\right)\left(1 - F\left(d^k\right)\right)v_i^{k-1}\left(s^{k-1}\right)\right.$$
$$\left. + \left(1 - F\left(d^k\right)\right)^2 v_i^{k-2}\left(s^{k-2}\right)\right) \cdot \frac{1}{1 - \left(F(d^k)\right)^2}$$
$$(8)$$

where $v_i^k = 0$ if $k \le 0$.

The original equation of the non-symmetric case (6) will be used later on to check whether a player has an incentive to deviate at E-state k from a given symmetric strategy combination.

The social optimum. In this section, we will study the supremum of the payoff of symmetric policies. Though this supremum may not be obtained, we will describe how to approach it, and find an approximation of it.

Since the policies we are studying are symmetric, we will denote by s^k the strategies of both physicians. Omitting the player index i, we will look for the supremum of $v^k(s^k)$ or, alternatively, $v^k(s^{k-1}, d^k)$. Let σ^k be the supremum of the payoff that can be achieved from E-state k onwards using a symmetric policy s^k:

$$\sigma^k = \sup_{s^k} v^k\left(s^k\right) = \sup_{s^{k-1}, d^k} v^k\left(s^{k-1}, d^k\right)$$

When we consider general policies, $d^k = 1$ is not allowed, because the payoff function cannot be defined at $d_i^k = d_{-i}^k = 1$, and the limit at this point does not exist. However, when we limit the discussion to symmetric policies, the definition of the payoff function can be extended to the case of $d^k = 1$.

Lemma 1.

$$\lim_{\epsilon \to 0} v^k\left(s^{k-1}, 1 - \epsilon\right) = \frac{k\alpha(r_2 - r_1)}{2} + v^{k-1}\left(s^{k-1}\right)$$

Proof See supplementary information (Supplementary Note 4)

By Lemma 1, we can extend the symmetric payoff function for $d^k = 1$, by defining

$$v^k\left(s^{k-1}, 1\right) = \frac{k\alpha(r_2 - r_1)}{2} + v^{k-1}\left(s^{k-1}\right)$$

with this definition, the strategy space becomes a closed set, and the payoff function is continuous in it. Therefore, the extended payoff function has a maximum value. Let $\left(\hat{d}^1, \hat{d}^2, \dots, \hat{d}^k, \dots\right)$ be a sequence of decision rules in the interval $[0, 1]$, such that $\hat{s}^k = \left(\hat{d}^1, \hat{d}^2, \dots, \hat{d}^k\right)$ maximizes v^k, that is, $\hat{\sigma}^k = v^k(\hat{s}^k)$. Note that if this maximum value is achieved by strategies containing some $d^k = 1$ then it cannot actually be obtained, but only approached.

The following theorem states the main characteristic of optimality in our model.

Theorem 2.

$$\lim_{k \to \infty} \hat{d}^k = 1$$

Proof See Supplementary Information (Supplementary Note 5).

The intuition behind this theorem is that since the model is characterized by endless patience: as long as a sufficient amount of antibiotic effectiveness still exists —it is worth waiting for the patients with a very high probability of infection (almost certainty).

The subgame perfect equilibrium. In this section, we analyze the model as a non-cooperative game between the two physicians. Our goal is to compare the socially optimal policy to the equilibrium strategy of the game, and in particular to Markov perfect equilibria (MPE) of the game. In order for it to be comparable to the optimal policy, we limit our discussion to pure symmetric equilibria of this type. In addition, based on our explanation of decision rules and strategies in section "Treatment policies in the dynamic model", when discussing the notion of equilibrium, we need only consider individual deviations to Markov strategies, and not to history-dependent strategies.

This section consists of three parts: in the section "The conditions for a symmetric pure-strategy Markov perfect equilibrium", we provide a criterion for a given policy to be an MPE, in the section "The existence of a Markov perfect equilibrium in pure symmetric strategies", we prove that such an equilibrium always exists and in the section "Comparing the MPE and the Social Optimum", we compare the MPE and the socially optimal policy.

The conditions for a symmetric pure-strategy Markov perfect equilibrium

The one-stage deviation principle. The analysis of subgame perfect equilibrium in our model is based on the one-stage deviation principle. This principle is well known in the game-theoretic literature; see for example Fudenberg and Tirole[59]: "in a finite multi-stage game with observed actions, strategy profile s is subgame perfect if and only if it satisfies the one-stage-deviation condition that no player i can gain by deviating from s in a single

stage and conforming to s thereafter" (p. 109). An extended definition and explanation of why this principle is valid in our model to appear in the Supplementary Information (Supplementary Note 6).

Based on the one-stage deviation principle, at every E-state, $k = 1, \dots, M$, the expected payoff under symmetric behavior (d^k, d^k) will be compared with the expected payoff of an individual stage-deviation from d^k to $\bar{d}^k \ne d^k$. A symmetric strategy profile s is an MPE if and only if

$$v_i^k\left(s^{k-1}, d^k, d^k\right) \ge v_i^k\left(s^{k-1}, \bar{d}^k, d^k\right) \text{ for every } k = 1, \dots, M \text{ and for all } 0 \le \bar{d} < 1 \quad (9)$$

and we will use backward induction to verify this.

Remark. If the condition for profile s is violated at E-state $k < M$, then any other strategy profile t with $t^k = s^k$ is not an MPE as well.

The condition (9) can be also represented as a set of difference-equations. The formulation of these equations will use the following definition:

For any given $0 \le d^k < 1$ let

$$g_{d^k}\left(\bar{d}^k\right) = \frac{\left(1 - \left[F(d^k)\right]^2\right)\int_{d^k}^1 p \cdot f(p)\mathrm{d}p - \left(1 - F(d^k)F(\bar{d}^k)\right)\int_{d^k}^1 p \cdot f(p)\mathrm{d}p}{(1 - F(d^k))^2\left(F(d^k) - F(\bar{d}^k)\right)}$$

The formulation of the equilibrium equations using $g_{d^k}\left(\bar{d}^k\right)$ appears in the Supplementary Information (Supplementary Note 7).

The following Lemma states an important property of $g_{d^k}\left(\bar{d}^k\right)$.

Lemma 3. $g_{d^k}\left(\bar{d}^k\right)$ is strictly increasing:

$$g_{d^k}\left(\bar{d}^k\right) > g_{d^k}\left(\bar{d}^k - \epsilon\right), \ \forall \ 0 \le \bar{d}^k < 1, \ 0 < \epsilon < \bar{d}$$

Proof See supplementary information (Supplementary Note 8).

Following the last part of this proof, we define

$$g_{d^k}\left(d^k\right) = \lim_{\bar{d}^k \to d^k} g_{d^k}\left(\bar{d}^k\right) = \frac{1 + F(d^k)}{1 - F(d^k)}d^k - \frac{F(d^k)\int_{d^k}^1 p \cdot f(p)\mathrm{d}p}{\left[1 - F(d^k)\right]^2} \quad (10)$$

This definition will serve us later on in proving the gap between the socially optimal policy and the MPE.

The unimodality of payoffs. The following theorem states that, when $k > 1$, for any given symmetric decision rule and assuming it is applied by the other physician, physician i faces a unimodal E-state payoff function with respect to his own decision rule. Furthermore, the E-state payoff function is strongly unimodal, i.e., the global maximum is attained at a single point, and the function is strictly increasing until that point and strictly decreasing thereafter.

Theorem 4. *It is impossible that the following two conditions hold simultaneously at any given E-state $k > 1$:*

$$v_i^k\left(s^{k-1}, d^k, d^k\right) \le v_i^k\left(s^{k-1}, d^k - x, d^k\right) \text{ for some } 0 < x \le d^k$$

and

$$v_i^k\left(s^{k-1}, d^k, d^k\right) \le v_i^k\left(s^{k-1}, d^k + y, d^k\right) \text{ for some } 0 < y < 1 - d^k$$

Proof See supplementary information (Supplementary Note 9).

This theorem has two important implications for proving the existence of a pure symmetric MPE:

First, when we want to check if a physician has an incentive to deviate from a given symmetric decision rule in a certain direction, it is enough to check a very small deviation in that direction (there are no local maxima or inflection intervals). We shall term this notion the minimal one-stage deviation principle. Second, if a physician has an incentive to deviate from a given symmetric decision rule in a certain direction, then he necessarily has no incentive to deviate in the other direction.

The existence of a Markov perfect equilibrium in pure symmetric strategies. Due to the one-stage deviation principle, in order to prove the existence of a symmetric pure MPE, we can use backward induction. Starting from E-state $k = 1$, we only need to prove that at every E-state k there exists a symmetric pure stage-equilibrium, $\left(d_1^k, d_2^k\right) = \left(\tilde{d}^k, \tilde{d}^k\right)$. Theorem 4 and the following lemma will enable us to do so.

Lemma 5.

$$\exists 0 < \epsilon, \lim_{d^k \to 1} v_i^k\left(s^{k-1}, d^k - \epsilon, d^k\right) > v_i^k\left(s^{k-1}, d^k, d^k\right)$$

Proof See Supplementary Information (Supplementary Note 10).

We can now prove the following existence theorem:

Theorem 6. *There always exists a symmetric pure-strategy MPE in the game.*

Proof We can prove the existence of a symmetric pure-strategy MPE by constructing it, using backward induction. At each E-state k, we will search for a symmetric pure-strategy stage-equilibrium $\left(\tilde{d}^k, \tilde{d}^k\right)$, given the symmetric pure-strategy equilibrium that was found in the previous E-states $\tilde{s}^{k-1}$. Due to the one-stage deviation principle, if these symmetric pure-strategy stage-equilibria exist for all $k = 1, \dots, M$ (i.e., no player i can gain by deviating from s in a single stage) then there exists a symmetric pure-strategy MPE in the game.

The following process is performed at each E-state $k = 1, \dots, M$:

First, we check whether $(0, 0)$ is a symmetric stage equilibrium. If it is—then an E-state-k-equilibrium exists $\left(\bar{d}^k, \tilde{d}^k\right) = (0, 0)$, and we can move to E-state $k + 1$.

Otherwise, the physicians must have an incentive to deviate upwards (to $\bar{d}_k > 0$). Let

$$B^k = \{\varepsilon | \exists \Delta > 0, \; v_i^k\left(s^{k-1}, \varepsilon + \Delta, \varepsilon\right) > v_i^k\left(s^{k-1}, \varepsilon, \varepsilon\right)\}$$

$$b^k = \sup B^k$$

$B^k \neq \varnothing$ because $(0, 0) \in B^k$.

$b^k \neq 0$ since the payoff function $v_i^k\left(s^{k-1}, \bar{d}_k, d^k\right)$ is continuous. Thus, if $v_i^k\left(s^{k-1}, 0 + \Delta, 0\right) > v_i^k\left(s^{k-1}, 0, 0\right)$ then there exist other symmetric combinations of decision rules $(d^k, d^k) = (\varepsilon, \varepsilon)$ close to $(0, 0)$ from which each physician has an incentive to deviate upwards.

$b^k \neq 1$ because by Lemma 5, if d^k is close enough to 1, then each physician has an incentive to deviate downwards, and by Theorem 4, if a physician has an incentive to deviate downwards from a given symmetric combination (d^k, d^k), then he does not have an incentive to deviate upwards.

We now check whether (b^k, b^k) is a symmetric stage equilibrium. If a player has an incentive to deviate upwards, then due to the continuity of the payoff function there exists ϵ such that each player also has an incentive to deviate upwards from $\left(b^k + \epsilon, b^k + \epsilon\right)$, but that contradicts the definition of b^k. Similarly, if a player has an incentive to deviate downwards, then due to the continuity of the payoff function there exists Δ and an interval from $(b^k - \Delta, b^k - \Delta)$ to (b^k, b^k) such that for every $0 < \epsilon < \Delta$, a player has an incentive to deviate downwards from $\left(b^k - \epsilon, b^k - \epsilon\right)$. But by Theorem 4, if at every $\left(b^k - \epsilon, b^k - \epsilon\right)$, a player has an incentive to deviate downwards then he does not have an incentive to deviate upwards, and that also contradicts the definition of b^k.

Therefore, $(\tilde{d}^k, \tilde{d}^k) = (b^k, b^k)$ is a symmetric pure-strategy stage equilibrium.

Comparing the MPE and the social optimum. After proving that a symmetric pure-strategy MPE always exists, we would like to compare it to the social optimum that was analyzed in "Reducing the problem" section. By Theorem 2, we know that as long as k is large enough the social optimum is waiting for patients with a very high probability of a bacterial infection (almost certainty). We now show that the MPE is significantly different, namely that in MPE the physicians always use antibiotics more extensively.

Theorem 7. *Let* $(\tilde{d}^k, \tilde{d}^k)$ *be a stage equilibrium.*

$$\lim_{k \to \infty} \tilde{d}^k \neq 1$$

Proof See supplementary information (Supplementary Note 11).

Reducing the problem. After proving the existing gap between the optimal policy and the equilibrium of the game and the players' rational incentive for overusing antibiotics, we strive to implement an approximation of the optimal policy as a new MPE of the game by coarsening the information available to the physicians. The motivation and practical interpretation of this process have been reviewed and explained in the results section. The following section contains its mathematical formulation and proofs of the main results.

Coarsening the information: calculations. We now replace the continuous information system $f(p)$ with a dichotomous discrete system signal system. We first need to determine a certain threshold probability T. The new system contains two signals, high (H) and low (L), indicating whether the patient's probability of having bacterial infection is higher or lower than the given threshold T. Each of these two signals appears in a certain probability and induces a different posterior that the patient has the infection (Table 4).

Payoffs and MPE conditions with a dichotomous signal system. In order to explore the effects of the information coarsening on our game, we first need to adjust our definitions of payoffs (Eqs. (6) and (8)) and MPE conditions (Eq. (S7.1)) to a discrete dichotomous information system.

Briefly, the full model contains recursive calculations of payoffs. $v_i^k\left(s^k\right)$ is the expected cumulative payoff of player i from E-state k onwards (from the current stage of antibiotic efficiency until the end of the game), for a strategy profile. It can be decomposed to $v_i^k\left(s^k\right) = v_i^k\left(s^{k-1}, d_i^k, d_{-i}^k\right)$, where d_i^k is the threshold of the posterior probability of a bacterial infection defining the decision rule of player i in E-state k. The full description of the model and the definitions of all the variables appear in subsections "Treatment policies in the dynamic model" and "Recursive calculation of payoffs" in the "Methods" section.

Using a dichotomous signal system is equivalent to limiting the set of available decision rules in each E-state to $d_i^k \in \{0, T\}$. That is, there are only a low and a high signal.

Thus, there are only four possible payoff combinations in each E-state: $v_i^k\left(s^{k-1}, 0, 0\right)$, $v_i^k\left(s^{k-1}, T, T\right)$, $v_i^k\left(s^{k-1}, 0, T\right)$, $v_i^k\left(s^{k-1}, T, 0\right)$.

For the extended calculations of the following equations see Supplementary Information (Supplementary Note 1).

If the strategy profile is symmetric, then the payoff of each player when they both treat everyone ($d^k = 0$) is:

$$v_i^k\left(s^{k-1}, 0, 0\right) = k\alpha(r_2 - r_1)(F(T)p_L + [1 - F(T)]p_H) + v_i^{k-2}\left(s^{k-2}\right) \quad (11)$$

and when both treat only patients with a high signal ($d^k = T$) it is:

$$v_i^k\left(s^{k-1}, T, T\right) = \left(k\alpha(r_2 - r_1)p_H + 2F(T)v_i^{k-1}\left(s^{k-1}\right) + (1 - F(T))v_i^{k-2}\left(s^{k-2}\right)\right) \cdot \frac{1}{1 + F(T)} \quad (12)$$

If the strategy profile is not symmetric, then the payoff of the physician who treats everyone is:

$$v_i^k\left(s^{k-1}, 0, T\right) = F(T)\left[k\alpha(r_2 - r_1)p_L + v_i^{k-1}\left(s^{k-1}\right)\right] + [1 - F(T)]\left[k\alpha(r_2 - r_1)p_H + v_i^{k-2}\left(s^{k-2}\right)\right] \quad (13)$$

and the payoff of the physician who treats only patients with a high signal is:

$$v_i^k\left(s^{k-1}, T, 0\right) = F(T)\left[v_i^{k-1}\left(s^{k-1}\right)\right] + [1 - F(T)]\left[k\alpha(r_2 - r_1)p_H + v_i^{k-2}\left(s^{k-2}\right)\right] \quad (14)$$

when we want to check whether a given symmetric strategy profile is an MPE, the one-stage deviation principle, wherein a strategy profile is an MPE if and only if no player can gain by deviating from it in a single stage and conforming to it thereafter[59], still applies (as in our original information setting, see "The conditions for a symmetric pure-strategy Markov perfect equilibrium" section). However, there are only two types of possible deviations—either from $d^k = T$ to $\bar{d}^k = 0$ or vice versa.

Thus, $d^k = 0$ (treating everyone) is a symmetric stage equilibrium if and only if there is no incentive to deviate to $d^k = T$ (treating only patients with a high signal). Formally, iff:

$$v_i^k\left(s^{k-1}, 0, 0\right) \geq v_i^k\left(s^{k-1}, T, 0\right)$$

Using (11) and (14) and some algebra we get

$$k\alpha(r_2 - r_1)p_L \geq v_i^{k-1}\left(s^{k-1}\right) - v_i^{k-2}\left(s^{k-2}\right) \quad (15)$$

and, similarly, $d^k = T$ (treating only patients with a high signal) is a symmetric stage equilibrium iff:

$$v_i^k\left(s^{k-1}, T, T\right) \geq v_i^k\left(s^{k-1}, 0, T\right)$$

Using (12) and (13) and some algebra we get

$$\frac{k\alpha(r_2 - r_1)([1 + F(T)]p_L - F(T)p_H)}{1 - F(T)} \leq v_i^{k-1}\left(s^{k-1}\right) - v_i^{k-2}\left(s^{k-2}\right) \quad (16)$$

Implementing optimal policy as a new MPE. When we use the new information system, the fixed symmetric policy of "treating only patients with a high signal" (i.e., $d^k = T$, $k = 1, \ldots, M$) can be considered a good approximation of the original optimal policy. Therefore, we would like to know whether it can be an MPE

In order to find the necessary and sufficient condition for this, we will use the following claim:

Claim 8. Under the fixed symmetric policy $s = (T, T)^M$ (i.e., $d^k = T$, $k = 1, \ldots, M$)

$$v_i^k\left(s^k\right) - v_i^{k-1}\left(s^{k-1}\right) > 0, \; k = 1, \ldots, M$$

That is, the fixed symmetric policy of "treating only patients with a high probability of a bacterial infection" yields a total expected payoff (from E-state k onwards) that is strictly increasing in k.

Proof See Supplementary Information (Supplementary Note 2)

Now we can use (16) and Claim 8 to set the following theorem.

Theorem 9. *Treating only patients with a high signal,* $s = ((T, T)^M)$, *is an MPE iff*

$$p_H - p_L \geq [1 - F(T)]p_H + F(T)p_L$$

Proof Due to the one-stage-deviation principle and Eq. (16) $s = ((T, T)^M)$ is an MPE if

$$\frac{k\alpha(r_2 - r_1)([1 + F(T)]p_L - F(T)p_H)}{1 - F(T)} \leq v_i^{k-1}\left(s^{k-1}\right) - v_i^{k-2}\left(s^{k-2}\right), \; k = 1, \ldots, M \quad (17)$$

Table 4 Information signals and their probability.

Signal	Probability of appearance	Posterior probability induced by the signal
H	$q_H = 1 - F(T) = \int_T^1 f(p)\,dp$	$p_H = \frac{\int_T^1 p \cdot f(p)\,dp}{\int_T^1 f(p)\,dp} = \frac{\int_T^1 p \cdot f(p)\,dp}{1 - F(T)}$
L	$q_L = F(T) = \int_0^T f(p)\,dp$	$p_L = \frac{\int_0^T p \cdot f(p)\,dp}{\int_0^T f(p)\,dp} = \frac{\int_0^T p \cdot f(p)\,dp}{F(T)}$

When $k = 1$:

$$v_i^{k-1}(s^{k-1}) - v_i^{k-2}(s^{k-2}) = 0$$

and by Claim 8, when $k > 1$:

$$v_i^{k-1}(s^{k-1}) - v_i^{k-2}(s^{k-2}) > 0$$

Therefore, if we want Eq. (17) to hold for all $k = 1, \dots, M$ we get

$$\frac{k\alpha(r_2 - r_1)([1 + F(T)]p_L - F(T)p_H)}{1 - F(T)} \leq 0$$

which means

$$[1 + F(T)]p_L \leq F(T)p_H$$

and finally

$$p_H - p_L \geq [1 - F(T)]p_H + F(T)p_L \qquad (18)$$

Data description and analysis. We obtained data of 1202 children aged <2 years who were hospitalized for bronchiolitis at Hillel Yaffe Medical Center in Hadera, Israel between 2008 and 2018. All children tested positive for RSV by antigen detection enzyme immunoassay, but 967 had only RSV bronchiolitis (viral infection), whereas 235 also had bacterial pneumonia, as confirmed in an X-ray scan. After retaining only variables with <10% missing values, our data contained 27 patient variables, including demographics (e.g., age, sex, place of birth), clinical symptoms and signs (e.g., temperature, tachypnea, etc.), comorbidities, and the season of hospitalization. Values were imputed using a random forest algorithm (using the randomForest R package[60]). We used the patient data to train a gradient boosted tree model (using the xgboost R package[61]) that classified bacterial and viral infections. Briefly, we tuned the hyperparameters of the model using a 10-fold cross validation, using the default values of the xgb.cv function in xgboost and a logistic loss function, and applied the model to the entire dataset to recover for each patient a value $0 < p < 1$. The estimated distribution of p, $\hat{f}(p)$, was a smoothed version of the resulting distribution of the patients, with a small positive constant (10^{-4}) added to the distribution to account for sampling limitations and create a support of $(0,1)$. This $\hat{f}(p)$ was considered as an approximation for the posterior distribution of having a bacterial infection. The final model output performed well with an area under the receiver operating curve (AUC) of 0.81. This study complied with all relevant ethical regulations for work with human participants. The study protocol was approved by the Institutional Review Board (Helsinki Committee) of Hillel Yaffe Medical Center. An exemption of informed consent was given by the Helsinki Committee given the retrospective study design. The patients' identities were kept confidential and coded information was used.

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

Data availability
Raw data were generated at Hillel Yaffe Medical Center. Derived data supporting the findings of this study are available upon reasonable request from the corresponding author.

Code availability
Code is available at https://github.com/uriobolski/GameTheoryAMR.

References
1. Bell, B. G., Schellevis, F., Stobberingh, E., Goossens, H. & Pringle, M. A systematic review and meta-analysis of the effects of antibiotic consumption on antibiotic resistance. *BMC Infect. Dis.* https://doi.org/10.1186/1471-2334-14-13 (2014).
2. Zaman, S. B. et al. A review on antibiotic resistance: alarm bells are ringing. *Cureus* **9**, e1403 (2017).
3. Dadgostar, P. Antimicrobial resistance: implications and costs. *Infect. Drug Resist.* **12**, 3903 (2019).
4. Bennett, J. E., Dolin, R. & Blaser, M. J. *Mandell, Douglas, and Bennett's Principles and Practice of Infectious Diseases* 2-Volume Set. 1 (Elsevier Health Sciences, 2014).
5. Poses, R. M., Cebul, R. D. & Wigton, R. S. You can lead a horse to water-improving physicians' knowledge of probabilities may not affect their decisions. *Med. Decis. Mak.* **15**, 65–75 (1995).
6. Denig, P. & Haaijer-Ruskamp, F. M. Therapeutic decision making of physicians. *Pharm. Weekbl.* **14**, 9–15 (1992).
7. Leibovici, L., Paul, M., Nielsen, A. D., Tacconelli, E. & Andreassen, S. The TREAT project: decision support and prediction using causal probabilistic networks. *Int. J. Antimicrob. Agents* **30**, 93–102 (2007).
8. Yamaga, S. & Shime, N. Association between appropriate empiric antimicrobial therapy and mortality from bloodstream infections in the intensive care unit. *J. Infect. Chemother.* **24**, 267–271 (2018).
9. Paul, M. et al. Systematic review and meta-analysis of the efficacy of appropriate empiric antibiotic therapy for sepsis. *Antimicrob. Agents Chemother.* **54**, 4851–4863 (2010).
10. Fraser, A. et al. Benefit of appropriate empirical antibiotic treatment: thirty-day mortality and duration of hospital stay. *Am. J. Med.* **119**, 970–976 (2006).
11. Goossens, H. et al. Outpatient antibiotic use in Europe and association with resistance: a cross-national database study. *Lancet* **365**, 579–587 (2005).
12. Crowther, G. S. & Wilcox, M. H. Antibiotic therapy and *Clostridium difficile* infection–primum non nocere–first do no harm. *Infect. Drug Resist.* **8**, 333 (2015).
13. Tedijanto, C., Olesen, S. W., Grad, Y. H. & Lipsitch, M. Estimating the proportion of bystander selection for antibiotic resistance among potentially pathogenic bacterial flora. *Proc. Natl. Acad. Sci. USA* **115**, E11988–E11995 (2018).
14. Tacconelli, E. Antimicrobial use: risk driver of multidrug resistant microorganisms in healthcare settings. *Curr. Opin. Infect. Dis.* **22**, 352–358 (2009).
15. Carrara, E., Pfeffer, I., Zusman, O., Leibovici, L. & Paul, M. Determinants of inappropriate empirical antibiotic treatment: systematic review and meta-analysis. *Int. J. Antimicrob. Agents* **51**, 548–553 (2018).
16. Lipsitch, M., Bergstrom, C. T. & Levin, B. R. The epidemiology of antibiotic resistance in hospitals: paradoxes and prescriptions. *Proc. Natl. Acad. Sci. USA* **97**, 1938–1943 (2000).
17. Obolski, U., Stein, G. Y. & Hadany, L. Antibiotic restriction might facilitate the emergence of multi-drug resistance. *PLoS Comput. Biol.* **11**, e1004340 (2015).
18. Obolski, U. & Hadany, L. Implications of stress-induced genetic variation for minimizing multidrug resistance in bacteria. *BMC Med.* **10**, 89 (2012).
19. Bergstrom, C. T., Lo, M. & Lipsitch, M. Ecological theory suggests that antimicrobial cycling will not reduce antimicrobial resistance in hospitals. *Proc. Natl. Acad. Sci. USA* **101**, 13285–13290 (2004).
20. Little, P. et al. Antibiotic prescription strategies for acute sore throat: a prospective observational cohort study. *Lancet Infect. Dis.* **14**, 213–219 (2014).
21. de la Poza Abad, M. et al. Prescription strategies in acute uncomplicated respiratory infections: a randomized clinical trial. *JAMA Intern. Med.* **176**, 21–29 (2016).
22. Jayashree, M. et al. Longitudinal comparative trial of antibiotic cycling and mixing on emergence of gram negative bacterial resistance in a pediatric medical intensive care unit. *J. Crit. Care* **56**, 243–248 (2020).
23. Naylor, N. R. et al. Estimating the burden of antimicrobial resistance: a systematic literature review. *Antimicrob. Resist. Infect. Control* **7**, 58 (2018).
24. Zhen, X., Lundborg, C. S., Sun, X., Hu, X. & Dong, H. Economic burden of antibiotic resistance in ESKAPE organisms: a systematic review. *Antimicrob. Resist. Infect. Control* **8**, 137 (2019).
25. Laxminarayan, R. & Brown, G. M. Economics of antibiotic resistance: a theory of optimal use. *J. Environ. Econ. Manage.* **42**, 183–206 (2001).
26. Qattan, W. T., Alkadi, K., Azmi, A. Aa. & Metwali, H. H. 19 Assessing physicians' compliance with medication-related clinical decision support alerts in the intensive care unit. *BMJ Open Quality* **8**, A8–A9 (2019).
27. Carroll, A. E., Anand, V. & Downs, S. M. Understanding why clinicians answer or ignore clinical decision support prompts. *Appl. Clin. Inform.* **3**, 309–317 (2012).
28. Kaplan, B. Evaluating informatics applications—clinical decision support systems literature review. *Int. J. Med. Inform.* **64**, 15–37 (2001).
29. Colman, A. M., Krockow, E. M., Chattoe-Brown, E. & Tarrant, C. Medical prescribing and antibiotic resistance: A game-theoretic analysis of a potentially catastrophic social dilemma. *PLoS ONE* **14**, e0215480 (2019).
30. Chen, X. & Fu, F. Social learning of prescribing behavior can promote population optimum of antibiotic use. *Front. Phys.* **6**, 139 (2018).
31. Maskin, E. & Sjostrom, T. in *Handbook of Social Choice and Welfare* (eds Arrow, K. J., Sen, A. K. & Suzumura, K.) 1, 237–288 (Elsevier, 2002).
32. Li, J. S. et al. Proposed modifications to the Duke criteria for the diagnosis of infective endocarditis. *Clin. Infect. Dis.* **30**, 633–638 (2000).
33. Wagner, F. P. & Mathiason, M. A. Using Centor criteria to diagnose streptococcal pharyngitis. *Nurse Pract.* **33**, 10–12 (2008).
34. Albrich, W. C., Monnet, D. L. & Harbarth, S. Antibiotic selection pressure and resistance in Streptococcus pneumoniae and Streptococcus pyogenes. *Emerg. Infect. Dis.* **10**, 514 (2004).
35. Yang, P. et al. Association between antibiotic consumption and the rate of carbapenem-resistant Gram-negative bacteria from China based on 153

tertiary hospitals data in 2014. *Antimicrob. Resist. Infect. Control* **7**, 1–7 (2018).

36. Rothwell, P. M. et al. A simple score (ABCD) to identify individuals at high early risk of stroke after transient ischaemic attack. *Lancet* **366**, 29–36 (2005).

37. Lloyd-Jones, D. M. et al. Framingham risk score and prediction of lifetime risk for coronary heart disease. *Am. J. Cardiol.* **94**, 20–24 (2004).

38. Callery, M. P., Pratt, W. B., Kent, T. S., Chaikof, E. L. & Vollmer, C. M. Jr. A prospectively validated clinical risk score accurately predicts pancreatic fistula after pancreatoduodenectomy. *J. Am. Coll. Surg.* **216**, 1–14 (2013).

39. Moxey, A. et al. Computerized clinical decision support for prescribing: provision does not guarantee uptake. *J. Am. Med. Informatics Assoc.* **17**, 25–33 (2010).

40. Khairat, S., Marc, D., Crosby, W. & Al Sanousi, A. Reasons for physicians not adopting clinical decision support systems: critical analysis. *JMIR Med. Informatics* **6**, e24 (2018).

41. Obolski, U., Dellus-Gur, E., Stein, G. Y. & Hadany, L. Antibiotic cross-resistance in the lab and resistance co-occurrence in the clinic: discrepancies and implications in *E. coli*. *Infect. Genet. Evol.* **40**, 155–161 (2016).

42. Paterson, D. L. "Collateral damage" from cephalosporin or quinolone antibiotic therapy. *Clin. Infect. Dis.* **38**, S341–S345 (2004).

43. Cherny, S. S. et al. Revealing antibiotic cross-resistance patterns in hospitalized patients through Bayesian network modelling. *J. Antimicrob. Chemother.* https://doi.org/10.1093/jac/dkaa408 (2020).

44. Morgan, D. J., Okeke, I. N., Laxminarayan, R., Perencevich, E. N. & Weisenberg, S. Non-prescription antimicrobial use worldwide: a systematic review. *Lancet Infect. Dis.* **11**, 692–701 (2011).

45. Hardin, G. The tragedy of the commons. *Science* **162**, 1243–1248 (1968).

46. Hiatt, H. H. Protecting the medical commons: who is responsible? *N. Engl. J. Med.* **293**, 235–241 (1975).

47. Weinstein, M. C. Should physicians be gatekeepers of medical resources? *J. Med. Ethics* **27**, 268–274 (2001).

48. Rajkomar, A., Dean, J. & Kohane, I. Machine learning in medicine. *N. Engl. J. Med.* **380**, 1347–1358 (2019).

49. Yelin, I. et al. Personal clinical history predicts antibiotic resistance of urinary tract infections. *Nat. Med.* **25**, 1143–1152 (2019).

50. Lewin-Epstein, O., Baruch, S., Hadany, L., Stein, G. & Obolski, U. Predicting antibiotic resistance in hospitalized patients by applying machine learning to electronic medical records. *Clin. Infect. Dis.* https://doi.org/10.1093/cid/ciaa1576 (2020).

51. Moran, E., Robinson, E., Green, C., Keeling, M. & Collyer, B. Towards personalized guidelines: using machine-learning algorithms to guide antimicrobial selection. *J. Antimicrob. Chemother.* **75**, 2677–2680 (2020).

52. Pauker, S. G. & Kassirer, J. P. The threshold approach to clinical decision making. *N. Engl. J. Med.* **302**, 1109–1117 (1980).

53. Pauker, S. G. & Kassirer, J. P. Therapeutic decision making: a cost-benefit analysis. *N. Engl. J. Med.* **293**, 229–234 (1975).

54. Von Neumann, J. & Morgenstern, O. *Theory of Games and Economic Behavior* (Princeton University Press, 1944).

55. Sadatsharifi, A., Davarpanah, M.-A., Namazi, S., Mottaghi, S. & Mahmoudi, L. Economic burden of inappropriate empiric antibiotic therapy: a report from southern Iran. *Risk Manag. Healthc. Policy* **12**, 339 (2019).

56. Andersson, D. I. & Hughes, D. Antibiotic resistance and its cost: is it possible to reverse resistance? *Nat. Rev. Microbiol.* **8**, 260 (2010).

57. Melnyk, A. H., Wong, A. & Kassen, R. The fitness costs of antibiotic resistance mutations. *Evol. Appl.* **8**, 273–283 (2015).

58. Puterman, M. L. *Markov Decision Processes: Discrete Stochastic Dynamic Programming* (John Wiley & Sons, 2014).

59. Fudenberg, D. & Tirole, J. *Game Theory* (MIT Press, 1991).

60. Liaw, A. et al. Classification and regression by randomForest. *R News* **2**, 18–22 (2002).

61. Chen, T., He, T., Benesty, M., Khotilovich, V. & Tang, Y. Xgboost: extreme gradient boosting. *R Packag. version 0.4-2* 1–4 (2015).

Author contributions

M.D., M.L., and D.S. designed and analyzed the original game theory approach; U.O. designed, performed, and analyzed the machine learning approach; S.B. and K.M. provided clinical insight and examples; K.M. and E.K. provided the clinical data; M.D. and U.O. wrote the manuscript, partially based on M.D.'s Ph.D. thesis, and all authors read and commented on the final version of the manuscript.

Competing interests

The authors declare no competing interests.

Additional information

Supplementary information The online version contains supplementary material available at https://doi.org/10.1038/s41467-021-21088-5.

Correspondence and requests for materials should be addressed to U.O.

Peer review information *Nature Communications* thanks Samit Bhattacharyya, Eunha Shim and Robert Vander Stichele for their contribution to the peer review of this work.

Reprints and permission information is available at http://www.nature.com/reprints

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