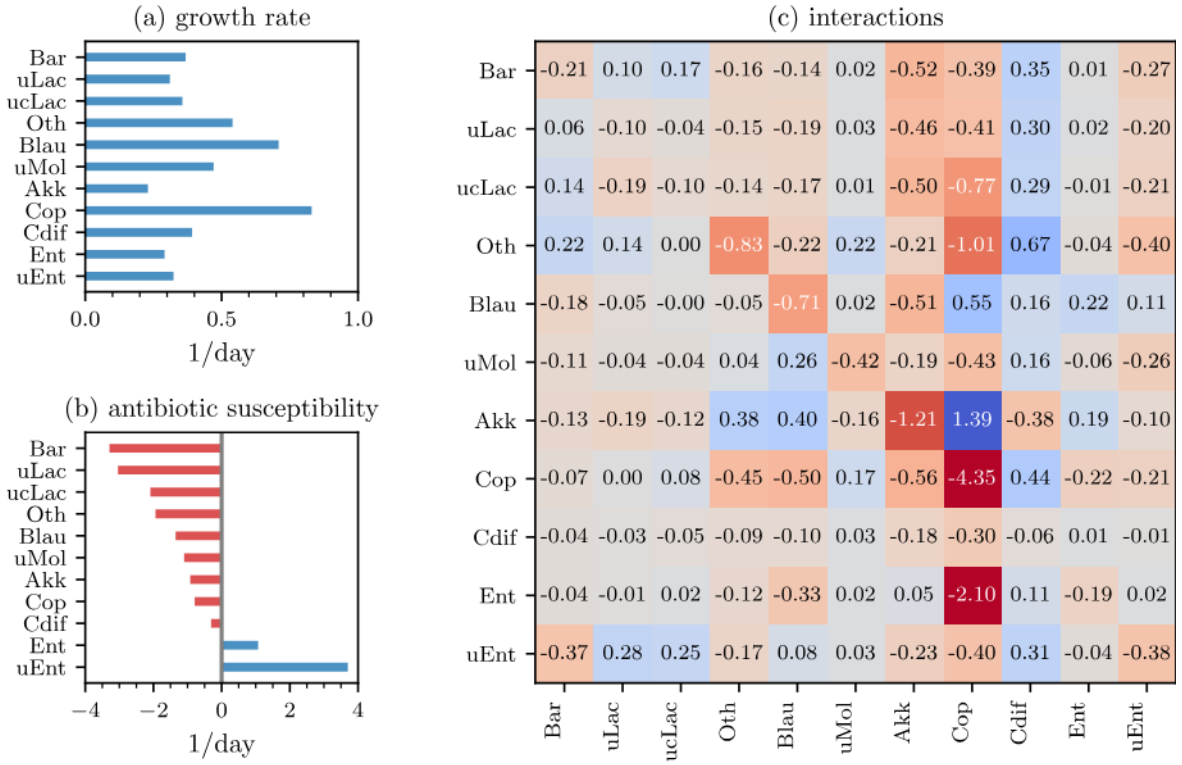

Fig. 10 Microbial community dynamics under antibiotic perturbation: (a) Species-specific growth rates (1/day). (b) Antibiotic susceptibility (1/day) in response to clindamycin, with negative values indicating inhibition and positive values indicating facilitation. (c) Interaction matrix M quantifying pairwise effects between species (element m_{ij} shows effect of species j on species i), with blue indicating repression and red indicating activation. Species are ordered by antibiotic susceptibility from most inhibited (bottom) to most promoted (top). Species abbreviations: Bar (*Barnesiella*), uLac (undefined Lachnospiraceae), ucLac (unclassified Lachnospiraceae), Oth (Other), Blau (*Blautia*), uMol (undefined Mollicutes), Akk (*Akkermansia*), Cop (*Coprobacillus*), Cdif (*C. difficile*), Ent (*Enterococcus*), uEnt (undefined Enterobacteriaceae). Data is based on [67, 68].

3.2 Population dynamics

Models of population dynamics are useful for studying species interactions in ecological and biomedical contexts. While we previously explored discrete-time formulations for multi-species predator–prey interactions (see Section 2.2), we now shift our focus to a continuous-time framework and consider the generalized Lotka–Volterra (gLV) equations. This framework has been employed, for instance, in microbial ecology [51], where species frequently engage in inhibitory and facilitative interactions. In this context, a commonly used form of the gLV equations is

$$\dot{x}_i(t) = x_i(t)\left(b_i + \sum_{j=1}^{n} m_{ij}x_j(t) + \epsilon_i u(t)\right), \tag{25}$$

where $x_i(t)$ denotes the abundance of microbial species i at time t, b_i is its intrinsic growth rate, and m_{ij} represents the interaction coefficient quantifying the effect of species j on species i. The term $\epsilon_i u(t)$ models the effect of an external antibiotic treatment $u(t)$, where ϵ_i indicates the susceptibility of species i. A positive ϵ_i corresponds to inhibition by the antibiotic, while a negative value allows for cases where species i may increase in abundance due to reduced competition following antibiotic exposure.

While gLV models are widely used to study microbial systems, a key limitation is their assumption of direct interactions between microbial species, which overlooks indirect effects mediated by competition for shared nutrients.

In [68], gLV parameters were inferred using mouse data from a study [67] that examined the effect of the antibiotic clindamycin on intestinal colonization by the spore-forming pathogen *C. difficile*. The dataset includes a total of $n = 11$ species. In Fig. 10, we show the estimated growth rates b_i, clindamycin susceptibilities ϵ_i, and elements of the interaction matrix m_{ij}. All growth rates are positive, while the diagonal elements of the interaction matrix, m_{ii}, are negative. These negative values indicate that each species can reach its carrying capacity even in the absence of other species. The inferred clindamycin susceptibilities suggest that the antibiotic inhibits all genera except *Enterococcus* and an

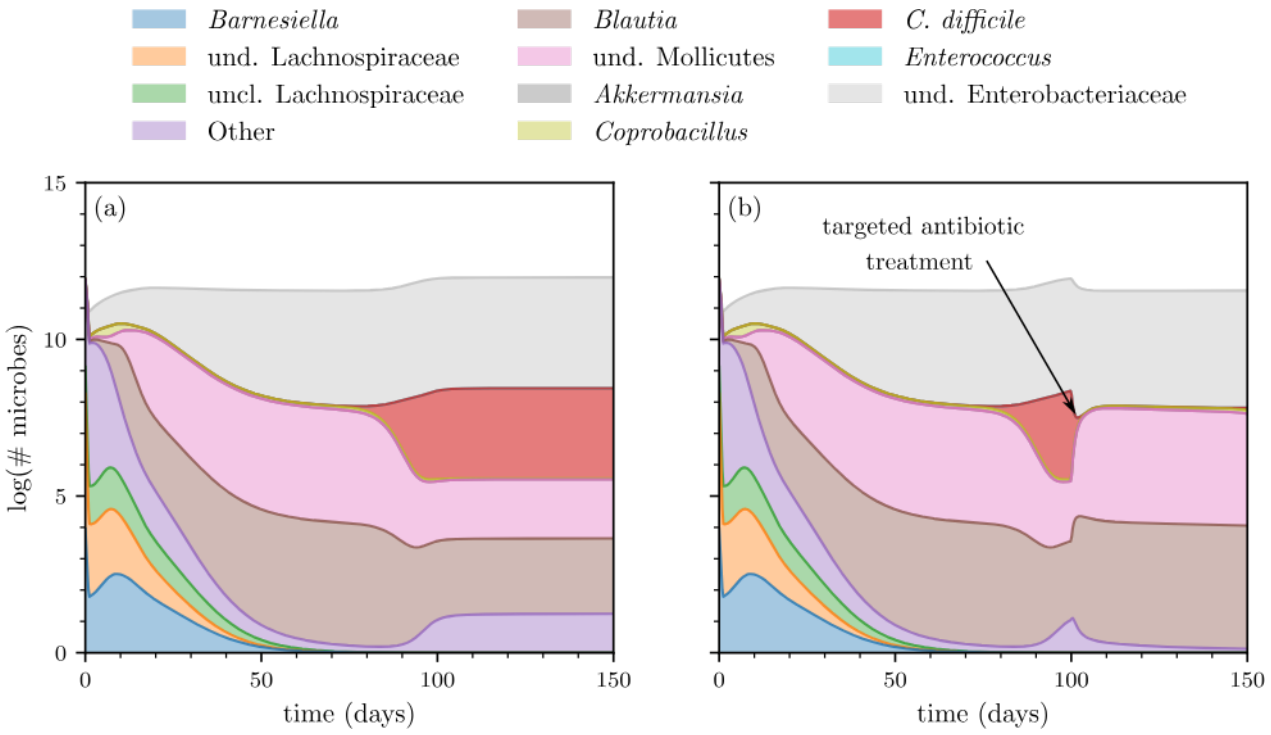

Fig. 11 Simulated microbial community dynamics under antibiotic interventions. (a) Without targeted treatment, the initial administration of clindamycin promotes the outgrowth of *C. difficile*, leading to a persistent infection. (b) A neural ODE controller applies a targeted antibiotic treatment starting on day 100 (black arrow), effectively suppressing *C. difficile*. Colors represent different microbial groups as indicated in the legend.

undefined group of *Enterobacteriaceae*. *C. difficile* itself appears to be only mildly inhibited by clindamycin.

We now consider a control problem focused on treating a *C. difficile* infection following the administration of clindamycin [69–71].

We use the growth rates, interaction coefficients, and clindamycin susceptibilities from [67,68] (see Fig. 10) and initialize the model with initial condition 5 from [68, 69]. To simulate infection onset, we introduce a small perturbation of 10^{-10} (in nondimensional units) to the *C. difficile* compartment and apply a unit dose of clindamycin on the first day. This treatment protocol is consistent with the constant dosing schedule considered in [69].

In Fig. 11(a), we show the evolution of microbial species. The results show that the initial antibiotic intervention, in combination with the *C. difficile* perturbation, leads to a substantial infection after approximately 90–100 days. In the absence of both perturbation and treatment, the system evolves as in Fig. 3(a) of [69].

C. difficile infections are, for instance, treated with antibiotics such as vancomycin or metronidazole [72]. Following the approach in [69], we now consider a hypothetical targeted antibiotic that is highly effective against *C. difficile*. The treatment begins on day 100 and lasts for 10 days. In our model, we set the antibiotic susceptibility of *C. difficile* to -1, while the susceptibilities of all other microbial species are set to zero. We train a neural ODE controller to minimize the loss function

$$\frac{1}{50}\int_{100}^{150} x_9(t)\,\mathrm{d}t + 0.01\frac{1}{10}\int_{100}^{110} \hat{u}(t;w)^2\,\mathrm{d}t\,, \tag{26}$$

which penalizes the abundance of *C. difficile* over time while also promoting prudent use of the targeted antibiotic.

In Fig. 11(b), we show the simulation results obtained using the trained controller. We observe that the targeted antibiotic treatment successfully suppresses the *C. difficile* infection. The neural ODE controller consists of five layers, each with four ELU-activated neurons. Training was performed for 200 epochs using the Adam optimizer with a learning rate of 10^{-3}, yielding a minimum loss of 0.061. As a baseline for comparison, we simulated a naive treatment strategy that administered a constant unit dose per day over the same 10-day period. This approach resulted in a loss more than ten times higher and failed to eliminate the infection.