Social norm dynamics in a behavioral epidemic model

Christos Charalambous$^{[0000-0002-0752-1007]}$

Department of Economics, University of Cyprus,
PO Box 20537, 1678 Nicosia, Cyprus
`charalambous.christos.2@ucy.ac.cy`

Abstract. Understanding the social determinants of preventive behaviour is crucial for epidemic modelling and policy design. We develop a behavioural epidemic model of seasonal disease on multilayer networks in which vaccination decisions combine experience-based learning with coevolving social norms. The framework distinguishes descriptive norms (what others do) from injunctive norms (what others believe ought to be done) and incorporates cognitive dissonance, social projection, and logical consistency. Simulations show that norm dynamics yield markedly different vaccination and infection outcomes than payoff-driven learning alone: injunctive norms exert stronger and more persistent effects than descriptive norms, and interventions targeting injunctive expectations are robust, whereas those acting on descriptive norms can be weaker or counterproductive. Once empirically validated, norm-based models can better capture human behaviour and inform strategies for collective-action problems beyond pandemics.

Keywords: Social norms, Behavioral epidemics, Agent-based modelling, Multiplex networks

1 Introduction

Vaccine hesitancy is among the most serious global public-health threats, underscored by COVID-19 [1] and the persistence of seasonal influenza. As a collective-action problem, it often resists purely regulatory solutions, motivating behavioural interventions that leverage social norms to promote cooperation and limit free-riding [2].

Behavioural–epidemic models capture feedbacks between disease dynamics and protective behaviour [3], while game-theoretic approaches model strategic vaccination under bounded rationality [4]. However, social influence is typically reduced to imitation or payoff averaging, neglecting psychological mechanisms—cognitive dissonance, social projection, and higher-order expectations—that shape belief formation and revision [2, 5]. This motivates a distinction between descriptive norms (what others do) and injunctive norms (what others believe should be done), whose alignment strongly enhances compliance and cooperation [2].

Empirical evidence shows that pro-vaccination norms increase uptake across diseases [6], and public-health strategies increasingly seek to exploit them [7].

Yet most models incorporate only descriptive norms via social learning [8], while injunctive norms—rooted in moral and reputational expectations—remain underexplored [9]. Few models integrate both dimensions [10], although recent psychologically grounded formulations provide a robust basis for doing so [5].

In contrast to imitation-based or static norm representations [4], we introduce a framework that distinguishes personal norms, empirical expectations, and normative expectations, linking their joint evolution to experience-weighted learning. This enables multiple layers of social norms to coevolve endogenously with epidemic risk, beyond standard payoff-based mechanisms. Building on cognitively grounded agent-based models of norm dynamics [10] and supported by experimental evidence on the superior effectiveness of injunctive-norm messaging [11], we develop an ABM that integrates a behavioural epidemic model with dynamic social norms [5] to study the coevolution of vaccination intentions and norms across repeated seasonal outbreaks.

We first present the model—its decision process, Experience-Weighted Attraction (EWA) learning [12], and norm dynamics—and then analyse how social norms shape vaccination and infection outcomes, including stubborn adherence and external interventions. Our contributions are fourfold: (i) coupling epidemic and norm dynamics through adaptive learning weights [12, 13]; (ii) incorporating cognitive dissonance, social projection, and logical consistency into norm evolution [9, 5]; (iii) demonstrating that injunctive norms exert stronger long-term effects on vaccination than descriptive norms, consistent with experimental evidence [14, 11]; and (iv) modelling interventions that act directly on norms rather than solely on information flows [15–17].

2 Model

During a pandemic, disease transmission and behavioural adaptation coevolve. We model infection dynamics with a Susceptible–Infected–Recovered (SIR) process on a physical contact network, while behavioural responses unfold on an overlapping social layer.

Both layers are empirically motivated synthetic networks. The physical layer follows a small-world topology [18] with mean degree 6 and rewiring probability 0.1, while the social layer is a Klimek–Thurner network [19, 20] (parameters $c = 0.58$, $r = 0.12$, $m = 1$), partially overlapping with the physical layer to reflect shared contact and information links. Results are robust to alternative topologies, including Erdos–Renyi networks on both layers. The two-layer architecture is shown in Fig. 1.

Each season, agents decide whether to vaccinate [21], one agent is initially infected, and the system evolves across seasons until equilibrium. Within each season, multiple SIR simulations are run using an event-driven algorithm [22] (with $\mu = 1$) to estimate infection probabilities. While social norms could in principle evolve continuously, we focus on endogenous norm formation driven by individual choices; accordingly, norms are updated only at the start of each season, when vaccination decisions are made. The full algorithm and schematic overview are provided in Fig. 1.

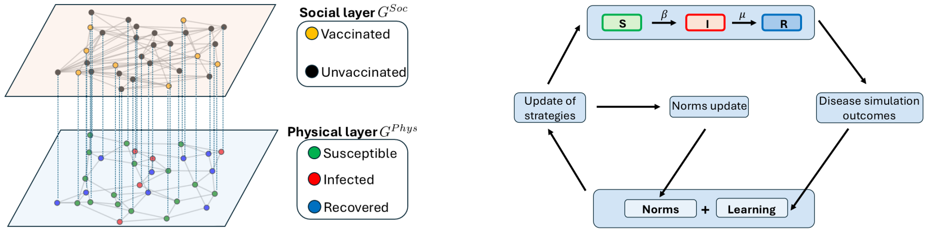

Fig. 1: **(a)** Vaccination game on a two-layer network: a physical layer hosting a seasonal SIR epidemic and a social layer governing vaccination decisions using information from both layers. **(b)** Schematic of the algorithm. Each season, SIR dynamics run on the physical layer to estimate infection risk; agents then decide whether to vaccinate based on payoff-based learning and normative factors, update their norms, and proceed to the next season.

To model behavioural adaptation, we use a game-theoretic framework in which vaccination payoffs depend on others' choices. Decisions are shaped not only by epidemiological risk but also by personal beliefs and social norms. Following [5], each individual i is characterized by three state variables: (i) a personal attitude y_i, representing beliefs about appropriate behaviour; (ii) an empirical expectation $\widetilde{x}_i$, capturing beliefs about others' intentions; and (iii) a normative expectation $\widetilde{y}_i$, reflecting beliefs about others' attitudes.

2.1 The decision making process

We define the dynamics of an agent's vaccination intention $x_i(t)$, i.e. the probability of choosing vaccination, which we assume to be driven by both material and normative considerations. Empirical evidence suggests that communities under risk rely more on experience and descriptive norms than on moral beliefs or injunctive norms [13]. In the absence of precise data on how perceived risk shapes this balance, we assume a simple linear relation between risk and the weight placed on empirical factors, following standard practice in norm–utility models [16]. Therefore we split the contribution to the intention $x_i\,(t+1)$ update as follows:

$$x_i\,(t+1) = (1 - S_i\,(t))\,x_i^{empirical}\,(t) + S_i\,(t)\,x_i^{injunctive}\,(t) \tag{1}$$

which depends on the safety parameter $S_i\,(t)$ that quantifies the safety an agent feels and takes values $0 < S_i\,(t) < 1$. The safety is assumed to be given by the following formula

$$S_i(t) = \begin{cases} 1 - \widehat{I}_i(t), & \text{if agent is unvaccinated,} \\ 1 - \widehat{I}_{i,\text{neighbors}}(t), & \text{if agent is vaccinated.} \end{cases} \tag{2}$$

where $\widehat{I}_i\,(t) = \frac{n_{sim}^{inf}(t)}{n_{sim}}$ with n_{sim} the total number of outbreak simulations and $n_{sim}^{inf}\,(t)$ the number of outbreak simulations when i got infected in the previous cycle, and $\widehat{I}_{i,neighbors}\left(t\right)=\frac{n_{sim}^{inf,neigh}\left(t\right)}{n_{sim}}$ with $n_{sim}^{inf,neigh}\left(t\right)$ the sum of the average number of infected neighbors over all simulation runs.

We define the contribution from what is or what agents think is actually happening, i.e. the empirical contribution to the intention, as

$$x_{i}^{empirical}\left(t\right)=\left(1-\phi_{i}\left(t\right)\right)p_{i}^{learn}\left(t\right)+\phi_{i}\left(t\right)\widetilde{x}_{i}\left(t\right) \tag{3}$$

where we see contributions from the agent's learning of material payoffs $p_{i}^{learn}\left(t\right)$ (see next subsection, $0\leq p_{i}^{learn}\left(t\right)\leq 1$) and what the agent thinks that the rest of the community actually chose on average $\widetilde{x}_{i}\left(t\right)$ (empirical expectations, $0\leq\widetilde{x}_{i}\left(t\right)\leq 1$). Their relative weight depends on $\phi_{i}\left(t\right)$, which determines whether decisions rely more on personal experience or on social information

$$\phi_{i}\left(t\right)=\sqrt{\phi_{i}^{change}\left(t\right)\phi_{i}^{consensus}\left(t\right)}. \tag{4}$$

$\phi_{i}^{change}\left(t\right)$ is the change-detector function and is defined as in [12]. $\phi_{i}^{consensus}\left(t\right)\in[0,1]$ is the consensus function given by

$$\phi_{i}^{consensus}\left(t\right)=\left|2*X_{i}-1\right| \tag{5}$$

where X_i is the average fraction of neighbors vaccinated as observed by the focal individual, and hence $\phi_{i}^{consensus}\left(t\right)$ is equal to 1 when neighbors agree and 0 if only half are vaccinated.

Intuitively, $\phi_{\text{change},i}$ measures how volatile neighbours' behaviour has been compared to the recent past, whereas $\phi_{i}^{\text{consensus}}(t)$, captures the degree of agreement among them. Using the geometric mean ensures that social information becomes influential only when both conditions are met—behaviour is stable and neighbours largely agree—whereas in noisy or polarized environments agents fall back on their own experience and attitudes.

The contribution from the belief dependent part of the agents, which is the mechanism activated when the agent is not at risk takes the following form

$$x_{i}^{injunctive}\left(t\right)=\left(1-\phi_{i}\left(t\right)\right)y_{i}\left(t\right)+\phi_{i}\left(t\right)\widetilde{y}_{i}\left(t\right) \tag{6}$$

where again the idea is to split the contributions between the effect of own beliefs and the beliefs of others using $\phi_i(t)$. We assume $0<y_i(t),\widetilde{y}_i(t)<1$. The dynamics of the norm variables $y_i(t)$ and $\widetilde{y}_i(t)$ are defined in the following subsections. A schematic of the contributions to total intention is shown in Fig. 2.

We adopt a minimal mechanism to model reliance on personal versus collective information. When individuals perceive instability or low consensus in neighbours' behaviour, they rely more on personal judgment, drawing on experience or internalized norms. This abstraction omits additional influences—such as trust in neighbours' decision-making under changing epidemiological conditions—which could depend on disease severity and would substantially increase model complexity.

Our approach builds on social and behavioural theories of adaptation [23] and the Continuous Opinions and Discrete Actions (CODA) framework [24], linking observed behaviour to evolving attitudes. It aligns with social-psychological models that emphasize cognitive rules over strict rationality in decision-making, and parallels recent applications in [16]. In contrast to earlier models treating attitudes as static, our framework captures feedback between attitudes and behaviour, enabling more realistic adoption dynamics and policy responses.

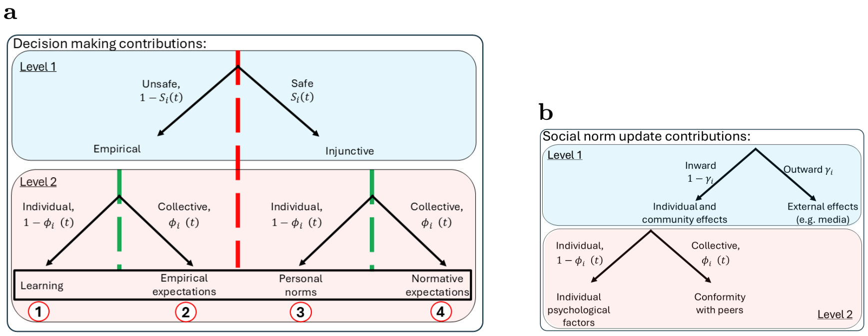

Fig. 2: **Decision-making process and social norm dynamics. a.: Schematic of the decision-making process.** Vaccination intention combines learning (1), personal norms (3), and social norms (2) and (4), each weighted by its relative influence. Agents rely more on empirical factors when they feel unsafe (low $S_i(t)$) or when their environment is unstable and heterogeneous, and more on normative factors under stable conditions. **b.: Schematic of social-norm dynamics**. Each agent updates norms based on external cues (with probability γ_i), personal beliefs, and peers' actions. In unstable and highly divided environments, agents rely more on internal factors than on social conformity.

2.2 The experience dependent mechanism

Human behaviour combines model-free and model-based reinforcement learning [12], integrating past experience with forward-looking reasoning. The Experience-Weighted Attraction (EWA) model captures both, unifying reinforcement and belief learning by weighting realized and forgone payoffs equally [10]. Forgone payoffs are inferred from neighbours' average infection rates, since individuals lack direct access to others' outcomes [25]. Agents evaluate payoffs over the last m cycles with a safety-dependent memory decay. EWA captures how individuals integrate past experience with expectations about others. This empirically validated approach [10] provides a cognitively grounded basis for modelling social adaptation and norm compliance.

The payoffs are therefore estimated as following:

$$\mathit{\Pi}_i^{Unvac}(t) = \begin{cases} 1 - c_I \widehat{I}_{i,neighbors}(t) & \mathit{if\ vaccinaed} \\ 1 - c_I \widehat{I}_i(t) & \mathit{if\ not\ vaccinated} \end{cases} \tag{7}$$

$$\mathit{\Pi}_i^{Vac}(t) = 1 - c_V$$

and based on this, the materially-motivated intention of the agent i at time t to vaccinate or not, is assumed to be given by the commonly employed Quantal Response Equilibrium (QRE) for binary choices, which is simply given by the logistic (softmax) function

$$p_i^{learn}(t) = \frac{1}{1 + e^{-\frac{\pi_i^{Unvac}(t) - \pi_i^{Vac}(t)}{\kappa}}} \tag{8}$$

Here $\pi_i^{Vac}(t)$ is the average payoff for the vaccinating option over the last m rounds, while $\pi_i^{Unvac}(t)$ is the average payoff received for not vaccinating over the last m rounds $\pi_i^{Unvac}(t)$, given by

$$\pi_i^{Unvac}(t) = \sum_{j=0}^{m} \frac{(S_i(t))^j}{\sum_{n=0}^{m}(S_i(t))^n} \mathit{\Pi}_i^{Unvac}(t-j) \tag{9}$$

where the safety parameter $S_i(t)$ defined above plays the role of a memory decaying function.

2.3 Social norm dynamics

To incorporate social norm dynamics into the EWA framework, we separate material payoffs from normative factors. Experimental evidence indicates that individuals represent others' beliefs and preferences independently of their own [14]. Accordingly, we model personal attitudes y_i, empirical expectations $\widetilde{x}_i$, and normative expectations $\widetilde{y}_i$ as distinct state variables, rather than reducing social influence to imitation, following [5].

Norm dynamics follow a DeGroot-type update driven by three psychological mechanisms: cognitive dissonance, aligning actions and self-beliefs; social projection, whereby individuals assume others are similar to themselves ; and logical consistency, reducing discrepancies between beliefs about others' actions and attitudes as in [5]. Attitudes and expectations thus adapt toward neighbours' average behaviour X_i [5] or externally promoted standards G_i^{ℓ}.

Rewriting the equations in [5], we obtain

$$\begin{aligned} y_i' &= y_i + \xi_i^1 \left[\widehat{C_i^{11}} x_i + \widehat{C_i^{12}} X_i + \widehat{C_i^{13}} G_i^1 - y_i\right] \\ \widetilde{y}_i' &= \widetilde{y}_i + \xi_i^2 \left[\widehat{C_i^{21}} y_i + \widehat{C_i^{22}} X_i + \widehat{C_i^{23}} G_i^2 - \widetilde{y}_i\right] \\ \widetilde{x}_i' &= \widetilde{x}_i + \xi_i^3 \left[\widehat{C_i^{31}} \widetilde{y}_i + \widehat{C_i^{32}} X_i + \widehat{C_i^{33}} G_i^3 - \widetilde{x}_i\right] \end{aligned} \tag{10}$$

where $\xi_i^{\ell} = \sum_j C_i^{\ell j}$ denotes the update rate of variable ℓ ($\ell = 1$ for y_i, $\ell = 2$ for $\widetilde{y}_i$, $\ell = 3$ for $\widetilde{x}_i$), and $C_i^{\ell j} = C_i^{\ell j} / \sum_k C_i^{\ell k}$ with $\sum_j C_i^{\ell j} = 1$. For simplicity, we set $\xi_i^1 = \xi_i^2 = \xi_i^3 = 1$. Results are robust across a wide parameter range, including hierarchies $\xi_i^1 < \xi_i^2 < \xi_i^3$, reflecting the intuition that personal norms adapt more slowly than empirical expectations. Future extensions could allow state-dependent ξ_i^{ℓ}, potentially generating qualitative transitions as in [26].

In this formulation, the coefficients $\widehat{C_i^{\ell j}}$ determine how each variable contributes to the evolution of the ℓth state. In particular, $\widehat{C_i^{\ell 1}}$ and $\widehat{C_i^{\ell 2}}$ represent self- and peer-weights: $\widehat{C_i^{\ell 1}}$ measures attraction toward the agent's own state, while $\widehat{C_i^{\ell 2}} = 1 - \widehat{C_i^{\ell 1}}$ captures neighbour influence. When present, $\widehat{C_i^{\ell 3}}$ quantifies sensitivity to externally promoted targets G_i^{ℓ}. In the baseline analysis we set $\widehat{C_i^{\ell 3}} = 0$, so norm evolution arises endogenously, and introduce $\widehat{C_i^{\ell 3}} > 0$ only when studying external interventions.

We assume $C_i^{\ell 1} = 1 - \phi_i(t), C_i^{\ell 2} = \phi_i(t)$, linking norm-updating weights to epidemic-driven change detection rather than fixing them exogenously. The function $\phi_i(t)$ combines change-detection $\phi_i^{\text{change}}(t)$ and consensus $\phi_i^{\text{consensus}}(t)$ terms, such that agents rely more on internal anchors under unstable conditions and more on peers when social signals are stable. This captures the idea that norm updating accelerates under clear consensus and slows under noise (Fig. 2B).

3 Results

For each season, we run 1000 simulations of a Susceptible–Infected–Recovered (SIR) process on networks of size $N = 500$, using an event-driven algorithm [22], with an independent network realization for each run.

We set the bounded-rationality parameter to $\kappa = 0.1$, corresponding to a quantal-response precision $\lambda = 1/\kappa = 10$. In quantal response theory (QRT), λ governs choice sensitivity, with higher values yielding more deterministic behaviour. The range $\lambda \in [1, 10]$ matches laboratory coordination and learning experiments [12], where higher-utility actions are chosen with 80–90% probability, making $\kappa = 0.1$ a realistic intermediate level of stochasticity commonly used in behavioural–epidemic models.

Vaccination and infection costs are set to $c_V = 1$ and $c_I = 0.1$, respectively, following standard assumptions in game-theoretic and norm-based vaccination models [9, 5] that infection is roughly an order of magnitude more costly than vaccination. Although not empirically calibrated, this ratio reproduces realistic voluntary vaccination levels observed in behavioural experiments [11] and reflects the widely held perception that infection entails substantially greater health and economic losses.

Unless otherwise stated, memory length is $m = 4$ and $\beta = 6$; results are qualitatively robust across wide parameter ranges. All figures report medians with lower and upper quartiles. Initial values of all three norm variables are drawn independently from a uniform distribution. In all cases, simulations are run for sufficiently many seasons to ensure convergence to evolutionary equilibrium, defined as a change in vaccination coverage of less than 0.025 over the final 50 seasons, with a maximum horizon of 200 seasons. All figures report median outcomes with lower and upper quartiles.

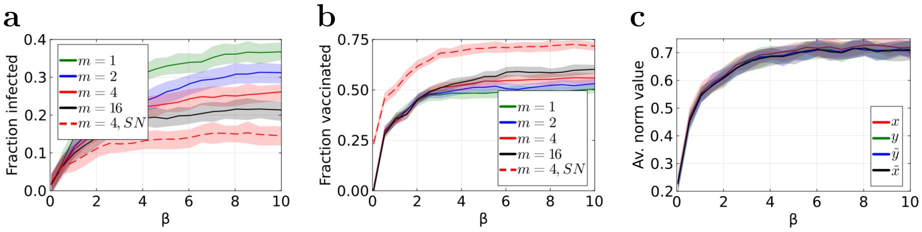

Fig. 3: **Role of memory and social norms in epidemic outcomes. a.** Infected fraction versus infectivity rate β. Infection increases with β and saturates at high values, while longer memory reduces outbreak size. For fixed $m = 4$, adding social norm dynamics further lowers infections. **b.** Vaccination coverage versus β. Coverage stabilizes near 0.5 for large β due to panic effects, but with norms, even small β sustain $\sim 20\%$ coverage as heterogeneous initial norms persist. **c.** Mean values of y, $\widetilde{y}$, $\widetilde{x}$, and x. All converge in the long run, indicating alignment between norms and vaccination intention.

The role of social norms We first examine memory effects on epidemic outcomes in the absence of social norms. Figure 3a shows that the infected fraction increases with infectivity β, remaining near zero for small β and saturating at high values, while longer memory reduces outbreak size by conditioning vaccination decisions on a richer payoff history. Figure 3b shows the corresponding vaccination coverage: for small β, coverage remains close to zero, whereas for large β it rises toward 0.5 due to panic-driven behaviour. Increasing memory further boosts coverage nonlinearly, reflecting network effects and the reduction of the safety factor $S_i(t)$ at higher β, which amplifies the influence of recent outcomes. Introducing social norm dynamics substantially alters infection levels and vaccination coverage. Norms sustain a non-trivial baseline vaccination rate ($\approx 20\%$) even at low β, reflecting persistent heterogeneity in initial attitudes. Figure 3c shows that the three norm variables converge at equilibrium and align with vaccination intentions.

The role of external factors We study the role of an external factor affecting the evolution of social norms, assuming that the external effect can be present independently of the epidemiological state of the community. This is introduced by setting $\widehat{C_i^{l3}} = \gamma_i^l$ where $0 \leq \gamma_i^l \leq 1$ for $l \in \{1,2,3\}$ is the strength of the external factor signal, and G_i^l is the target value for the norms. $\widetilde{C_i^{lj}} = \left(1-\gamma_i^l\right)\widehat{C_i^{lj}}$, where $\widehat{C_i^{l1}}$ and $\widehat{C_i^{l2}} = 1 - \widehat{C_i^{l1}}$ as before. Note that the condition that the coefficients of the variables sum up to 1 is still satisfied such that these can be still interpreted as relative contributions. In our studies, we assume the same strength γ_i^l and target G_i^l for all agents, i.e. $\gamma_i^l = \gamma^l$ and $G_i^l = G^l$ $\forall i \in \{1,...,N\}$.

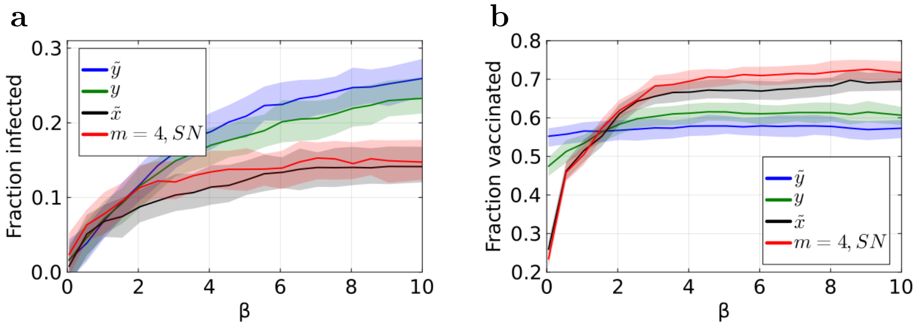

Fig. 4: **Role of external factors on social norms. a.** Infected fraction versus infectivity rate β when an external factor influences y_i, $\widetilde{y}_i$, or $\widetilde{x}_i$. The factor reduces perceived disease severity, driving norms toward $G_i = 0.6$, below the equilibrium value 0.7 in Fig. 3. It increases infections when applied to y_i or $\widetilde{y}_i$, but has negligible effect on $\widetilde{x}_i$. **b.** Vaccination coverage versus β for the same cases. Parameters: $\gamma_i = 1.0$, $G_i = 0.6$.

Figures 4 and 5 examine how external influences on social norms affect epidemic outcomes. Figure 4 shows that when an external factor downplays disease severity by shifting norms toward a lower target ($G = 0.6$), outbreaks increase if the intervention acts on personal norms or injunctive expectations, whereas interventions on empirical expectations have negligible effects.

Figure 5 reveals richer dynamics. Panel a shows that outbreak size depends jointly on the targeted norm, intervention strength γ, and target value G. Driving personal norms or injunctive expectations toward lower values consistently increases outbreaks, while interventions on empirical expectations can induce resistance: for intermediate targets ($G \approx 0.6$–0.7), infection levels fall below the no-intervention baseline, indicating resilience to external influence. Panel b highlights these nonlinearities further. For targets near $G \approx 0.7$, outbreaks can shrink under both injunctive and empirical expectation interventions, with empirical expectations providing the strongest buffering effect, whereas for low targets ($G < 0.6$) outbreaks increase as expected. For intermediate values ($G \approx 0.4$–0.6), interventions on personal or injunctive norms remain more effective, while for sufficiently low G targeting empirical expectations may dominate.

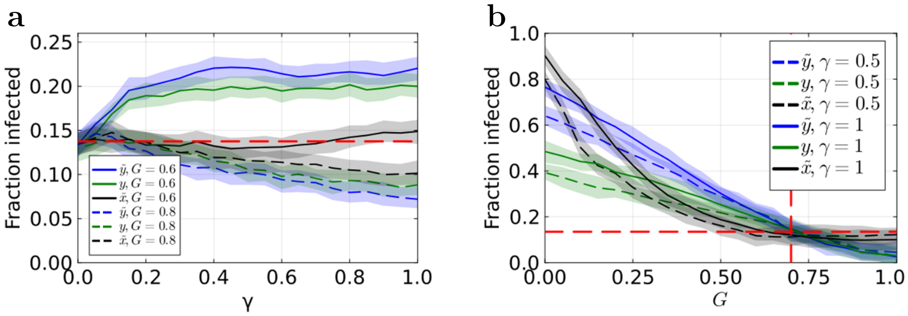

Fig. 5: **Role of external strength γ and target value G. (a)** Outbreak size as a function of coupling strength γ when the external factor acts on personal attitudes y_i, normative expectations $\widetilde{y}_i$, or empirical expectations $\widetilde{x}_i$, for targets $G = 0.6$ (lower perceived severity) and $G = 0.8$ (higher). The red dashed horizontal line denotes the baseline infection level without intervention. Driving y_i or $\widetilde{y}_i$ toward $G = 0.6$ increases infections, whereas applying the same force to $\widetilde{x}_i$ has little effect. For $G = 0.8$, outbreaks decrease in all cases, with the strongest reduction when targeting $\widetilde{y}_i$ and the weakest when targeting $\widetilde{x}_i$. **(b)** Outbreak size as a function of G for $\gamma = 0.5$ and $\gamma = 1$. The red dashed horizontal line marks the baseline infection level, and the vertical line indicates the corresponding vaccination level. For large G (nudging toward vaccination), interventions on moral-related norms (y_i or $\widetilde{y}_i$) are highly effective, driving infections toward zero, whereas interventions on $\widetilde{x}_i$ remain substantially less effective. For small G, outbreaks increase. Notably, for intermediate values $G \approx 0.6$–0.7, interventions on $\widetilde{x}_i$ can buffer external influence (e.g., $\gamma = 0.5$, black dashed line), yielding outbreak sizes below the no-intervention baseline.

Policy relevance Distinguishing injunctive from descriptive norm interventions is essential for effective public-health messaging. Injunctive messages convey socially or morally expected behaviour (e.g., appeals from trusted authorities), whereas descriptive messages report prevailing behaviour (e.g., vaccination rates). Our simulations show that descriptive messages can backfire when they highlight uptake below the counterfactual baseline (e.g., $G = 0.6$, $\gamma = 1$ for baseline uptake 0.7), inadvertently signalling that non-vaccination is common and reducing coverage.

Consistent with empirical evidence [13, 14], interventions targeting perceived approval and moral norms have stronger effects on vaccination intentions, particularly under uncertainty about others' behaviour. Our model formalizes how such cues, combined with probabilistic decision-making and social learning, shape collective outcomes. In contrast to multilayer approaches emphasizing structural information–disease feedbacks [3], we focus on behavioural feedbacks driven by evolving normative expectations and adaptive learning, revealing a clear asymmetry: interventions on injunctive norms propagate system-wide, whereas those targeting empirical expectations can induce resilience or neutralize attempts to downplay risk, in line with experimental [14] and theoretical evidence [27].

The simulations highlight three general principles of norm–epidemic coevolution: (i) even simple norm dynamics can suppress outbreaks by sustaining vaccination at low infectivity and reducing behavioural volatility beyond imitation or reinforcement alone [4]; (ii) norm components are asymmetric, with interventions on normative expectations exerting stronger collective effects than those on empirical expectations, consistent with the coordinating role of perceived approval [11, 14] for sufficiently high target values G; and (iii) norm interventions can generate nonlinear and counterintuitive effects [1, 13]. Overall, these results show how the structure of norm dynamics shapes epidemic and behavioural outcomes, complementing earlier analyses of awareness–disease coupling in multilayer networks [3].

4 Discussion

Most epidemic–behavior models rely on imitation or payoff-based reinforcement, overlooking how attitudes, expectations, and social pressures shape preventive decisions. We address this gap with a behavioral epidemic model that couples disease spread with Experience-Weighted Attraction (EWA) learning [12], integrating reinforcement and belief learning via realized and forgone payoffs. The framework embeds social norm dynamics distinguishing descriptive and injunctive norms [2, 5], and incorporates cognitive dissonance, social projection, and consistency mechanisms governing alignment with peers and authorities.

We show that coevolving norms strongly affect outbreak size and vaccination coverage. Interventions targeting personal norms or injunctive expectations have robust effects, whereas those acting on descriptive norms are weaker or can be counterproductive. These results align with experimental evidence: in the "transmission game," injunctive-norm messages reduced risk-taking more effectively than descriptive information or case counts [11]. This correspondence suggests a route for empirical calibration via longitudinal experiments or panel surveys tracking the joint evolution of personal norms, empirical expectations, and perceived approval under repeated interventions [1], and may help explain the asymmetric recovery of social norms reported by Vriens *et al.* [26].

The model is intentionally stylized and explanatory rather than predictive, isolating core mechanisms underlying the coevolution of epidemic risk and social norms. Assumptions of homogeneous agents, fixed social networks, and stable information environments enable controlled analysis but abstract from media and institutional feedbacks. Complementary work incorporating regret and uncertainty [28] explores additional emotional and informational channels. Future extensions could introduce agent heterogeneity, adaptive or multilayer networks [3, 20], and empirical calibration using experimental or survey data.

More broadly, empirically grounded norm-based models can inform communication strategies and interventions across collective-action domains, including vaccination, misinformation, and environmental cooperation [6, 7]. Further extensions could incorporate homophily in intentions or vaccination status, as well as peer sanctions and rewards in norm evolution.

Acknowledgments. I am thankful to Prof. Sandro Meloni for insightful discussions. This project received funding from the European Union's Horizon 2020 research and innovation program under the Marie Sklodowska-Curie Grant Agreement No 101034403.

Data Availability Statement Replication data and code can be found in GitHub: `https://github.com/Christos3788/Social-norm-dynamics-in-a-Behavioral-epidemic-model`.

References

1. Vriens, E., Tummolini, L. & Andrighetto, G. Vaccine-hesitant people misperceive the social norm of vaccination. *PNAS Nexus* **2**, pgad132; 10.1093/pnasnexus/pgad132 (2023).
2. Bicchieri, C. *The Grammar of Society: The Nature and Dynamics of Social Norms.* (Cambridge University Press, Cambridge, 2005).
3. Wang, Z., Andrews, M. A., Wu, Z. X., Wang, L. & Bauch, C. T. Coupled disease–behavior dynamics on complex networks: a review. *Phys. Life Rev* **15**, 1–29; 10.1016/j.plrev.2015.07.006 (2015).
4. Reluga, T. C. Game theory of social distancing in response to an epidemic. *PLOS Comput. Biol.* **6**, e1000793; 10.1371/journal.pcbi.1000793 (2010).
5. Gavrilets, S. Coevolution of actions, personal norms, and beliefs about others. *Evol. Hum. Sci.* **3**, e44; 10.1017/ehs.2021.40 (2021).
6. Moehring, A. *et al.* Providing normative information increases intentions to accept a COVID-19 vaccine. *Nat. Commun.* **14**, 126; 10.1038/s41467-022-35052-4 (2023).
7. Brewer, N. T., Chapman, G. B., Rothman, A. J., Leask, J. & Kempe, A. Increasing vaccination: putting psychological science into action. *Psychol. Sci. Public Interest* **18**, 149–207; 10.1177/1529100618760521 (2017).
8. Bauch, C. T. & Bhattacharyya, S. Evolutionary game theory and social learning can determine how vaccine scares unfold. *PLOS Comput. Biol.* **8**, e1002452; 10.1371/journal.pcbi.1002452 (2012).
9. Oraby, T., Thampi, V. & Bauch, C. T. The influence of social norms on the dynamics of vaccinating behaviour for paediatric infectious diseases. *Proc. R. Soc. B* **281**, 20133172; 10.1098/rspb.2013.3172 (2014).
10. Realpe-Gómez, J., Andrighetto, G., Nardin, L. G. & Montoya, J. A. Balancing selfishness and norm conformity can explain human behavior in large-scale prisoner's dilemma games and can poise human groups near criticality. *Phys. Rev. E* **97**, 042321; 10.1103/PhysRevE.97.042321 (2018).

11. Woike, J. K., Hafenbrädl, S., Kanngiesser, P. & Hertwig, R. The transmission game: testing behavioral interventions in a pandemic-like simulation. *Sci. Adv.* **8**, eabk0428; 10.1126/sciadv.abk0428 (2022).
12. Ho, T.-H., Camerer, C. F. & Chong, J.-K. Self-tuning experience weighted attraction learning in games. *J. Econ. Theory* **133**, 177–198; 10.1016/j.jet.2005.12.008 (2007).
13. Heiman, S. L., Garrett, N., Vosgerau, J., Tinghög, G. & Västfjäll, D. Descriptive norms caused increases in mask wearing during the COVID-19 pandemic. *Sci. Rep.* **13**, 11856; 10.1038/s41598-023-38593-w (2023).
14. Szekely, A., Lipari, F., Antonioni, A., Paolucci, M., Sánchez, A., Perc, M. & Andrighetto, G. Evidence from a long-term experiment that collective risks change social norms and promote cooperation. *Nat. Commun.* **12**, 5452; 10.1038/s41467-021-25734-w (2021).
15. Meiske, B., Álvarez Benjumea, A., Andrighetto, G. & Polizzi, E. Nudging punishment against sharing of fake news. *Eur. Econ. Rev.* **168**, 104795; 10.1016/j.euroecorev.2024.104795 (2024).
16. Tverskoi, D., Babu, S. & Gavrilets, S. The spread of technological innovations: effects of psychology, culture and policy interventions. *R. Soc. Open Sci.* **9**, 211833; 10.1098/rsos.211833 (2022).
17. Gavrilets, S. & Richerson, P. J. Authority matters: propaganda and the coevolution of behaviour and attitudes. *Evol. Hum. Sci.* **4**, e51; 10.1017/ehs.2022.48 (2022).
18. Salathé, M., Kazandjieva, M., Lee, J. W., Levis, P., Feldman, M. W. & Jones, J. H. A high-resolution human contact network for infectious disease transmission. *Proc. Natl. Acad. Sci. USA* **107**, 22020–22025; 10.1073/pnas.1009094108 (2010).
19. Klimek, P. & Thurner, S. Triadic closure dynamics drives scaling laws in social multiplex networks. *New J. Phys.* **15**, 063008; 10.1088/1367-2630/15/6/063008 (2013).
20. Charalambous, C., Sanchez, D. & Toral, R. Language competition on coevolving networks: an agent-based approach of nodes and links coevolution. *Front. Complex Syst.* **1**, 1304448; 10.3389/fcpxs.2023.1304448 (2023).
21. Cardillo, A., Reyes-Suárez, C., Naranjo, F. & Gómez-Gardeñes, J. Evolutionary vaccination dilemma in complex networks. *Phys. Rev. E* **88**, 032803; 10.1103/PhysRevE.88.032803 (2013).
22. Kiss, I. Z., Miller, J. C. & Simon, P. L. *Mathematics of epidemics on networks.* Springer, Cham; 10.1007/978-3-319-50806-1 (2017).
23. Davis, R., Campbell, R., Hildon, Z., Hobbs, L. & Michie, S. Theories of behaviour and behaviour change across the social and behavioural sciences: a scoping review. *Health Psychol. Rev.* **9**, 323–344; 10.1080/17437199.2014.941722 (2015).
24. Martins, A. C. R. Continuous opinions and discrete actions in opinion dynamics problems. *Int. J. Mod. Phys. C* **19**, 617–624; 10.1142/S0129183108012339 (2008).
25. Sánchez, A. Physics of human cooperation: experimental evidence and theoretical models. *J. Stat. Mech.* **2018**, 024001; 10.1088/1742-5468/aaa388 (2018).
26. Vriens, E., Andrighetto, G. & Tummolini, L. Risk, uncertainty, & social norms in vaccination decisions. *Philos. Trans. R. Soc. B* **379**, 20230035; 10.1098/rstb.2023.0035 (2024).
27. Roozmand, O., Deffuant, G., Andrighetto, G. & Paolucci, M. Simulating collective risk management from experimental data. In *Social Simulation Conference 2022* (2022).
28. Charalambous, C. Regret, Uncertainty, and Bounded Rationality in Norm-Driven Decisions. arXiv:2511.10342 (2025).