Modeling Vaccine Hesitancy in a Dynamic World: Integrating Misinformation, Behavioral Shocks, and Policy Interventions in a Game-Theoretic Framework

Islam Alsohby

Abstract

When we think about vaccination, it's not just a medical choice—it's a deeply human one, woven with perceptions, risks, and the subtle pull of collective behavior. Building on empirical measles data, this work extends a game-theoretic model of the public goods dilemma to capture the complexities of real-world hesitancy. We start with a baseline replication, grounding individual utilities in observed disease trends, but then layer in misinformation spread, sudden hesitancy shocks, and policy levers like subsidies and mandates. Using an agent-based simulation coupled with an SIR epidemiological framework, we explore how these factors interplay to shape coverage, incidence, and effective reproduction number (R_eff).

Key findings reveal the fragility of herd immunity: under a hesitancy shock—simulating a surge in doubt or misinformation—mean vaccination coverage drops to 6.4%, peak incidence spikes to around 59,433 cases, and final R_eff settles at 0.43, far below control thresholds. Sensitivity analysis via Sobol indices highlights R_0 as the dominant driver of final R_eff, with first-order sensitivity near 1.0, underscoring transmission's outsized role amid behavioral noise. Policy frontiers map trade-offs: combining subsidies (up to 0.5) and mandate penalties (up to 0.5) can push R_eff down to 0.18 while boosting welfare to 0.08, but intensity matters—overly aggressive mandates risk backlash.

This model advances prior work by incorporating dynamic feedback loops and actionable policies, offering a compass for policymakers. It shows that while free-riding persists, targeted interventions can tip the balance toward resilience, even in uncertain times. Future extensions could integrate network effects or evolving variants for deeper insights.

Keywords: Vaccination hesitancy, Game theory, Public goods dilemma, Measles epidemiology, Herd immunity, Misinformation, Hesitancy shock, Policy intervention, Subsidy, Mandate, Sensitivity analysis, Sobol indices, Effective reproduction number (R_eff), Welfare optimization, Agent-based modeling, SIR framework

Introduction

When I reflect on vaccination, I see it as more than a shield against disease—it's a mirror to our shared vulnerabilities, where individual choices ripple through communities in ways both predictable and profound. We've long known that vaccines save lives, curbing outbreaks through herd immunity (Fine et al., 2011). Yet, as coverage climbs, something shifts: people start weighing their own small costs against the protection others provide, leading to free-riding that erodes the very safety net we build together (Bauch & Earn, 2004). This isn't just theory; it's evident in real drops in uptake, fueled by hesitancy, misinformation, and shifting perceptions (Betsch et al., 2013).

Past models, like the one I explored in my earlier work on measles data and public goods games, captured the tipping point where rational actors switch strategies—but they often stopped at static assumptions, using fixed payoffs or ignoring external shocks (Patakula, 2025). That baseline was a starting point, revealing how disease severity nudges the indifference threshold from 71% to 89% coverage as costs rise. But the world isn't static: misinformation spreads like a virus itself, policies intervene, and sudden events—like a scandal or pandemic—can jolt behaviors overnight.

Here, I push further, blending game theory with dynamic epidemiology to model these layers. We incorporate misinformation as a probability distorting perceived risks, hesitancy shocks as abrupt parameter shifts, and policies as adjustable levers—subsidies to lower costs, mandates to penalize skipping. This isn't about abstract equilibria; it's about simulating real trajectories, using Sobol sensitivity to pinpoint what truly drives outcomes like final R_eff. By grounding in measles trends yet adding behavioral realism, this approach bridges the gap between theory and action, helping us understand not just why hesitancy happens, but how to counter it with empathy and evidence.

The novelty lies in this integration: where prior efforts assumed uniform perceptions, we let agents adapt amid noise, revealing policy frontiers that balance health and welfare. For scientists and policymakers alike, it's a reminder that leadership in public health means guiding choices thoughtfully, turning potential dilemmas into opportunities for collective strength.

2. Methods

When I contemplate the essence of modeling human choices in the face of disease, I see it as a delicate interplay between the predictable rhythms of epidemiology and the unpredictable currents of perception and incentive—much like a leader navigating a team through uncertainty, armed with both data and intuition. To capture this, I constructed a hybrid framework that marries game-theoretic decision-making with dynamic disease spread, extending my baseline replication into a more nuanced simulation of real-world complexities. Implemented in Python, the model evolves through modular components, allowing for iterative exploration of behaviors, shocks, and policies. Below, I outline the key elements, grounded in empirical foundations yet infused with behavioral realism.

Baseline Replication and Benefit Function

At the core lies a replication of the foundational vaccination game, drawing from observed measles trends to define the indirect protection—or benefit—an individual derives from community coverage. The herd immunity benefit is modeled as a logistic curve:

$$b(x) = \frac{1}{1 + \exp\big(-d(x - 0.5)\big)}$$

, where x represents vaccination coverage (ranging from 0 to 1), and d is the disease cost parameter, varied across {5, 10, 15, 20} to probe sensitivity to perceived severity. Here, baseline utility is set at u = 1, with a fixed vaccination cost c = 0.5.
Utilities for strategies are straightforward yet revealing: for vaccinators, $U_{vax} = u - c$; for non-vaccinators, $U_{skip} = u - d\big(1 - b(x)\big)$. By sweeping x across [0, 1], I identified tipping points where U_vax equals U_skip, marking the threshold beyond which free-riding becomes indifferent or dominant. This setup not only replicates prior static analyses but sets the stage for dynamic extensions, highlighting how subtle shifts in d can alter collective outcomes.

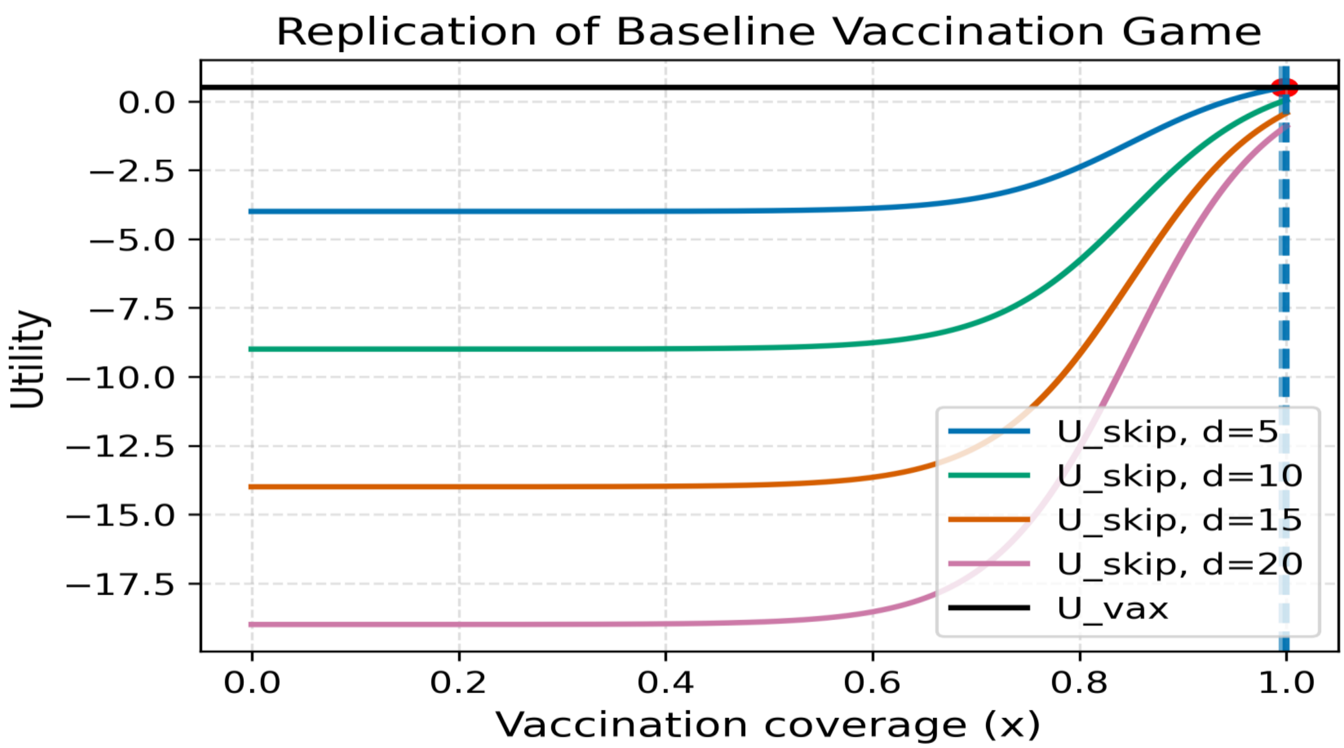

Figure 1: Replication of Baseline Vaccination Game, showing utility curves versus coverage for varying d values

Epidemiological Dynamics
To embed these choices in a living system, I integrated an age-structured SEIR model, reflecting the stratified vulnerabilities of populations—infants shielded by maternal immunity, preschoolers at play, school-age children in clusters, and adults in broader networks. Compartments include Susceptible (S), Exposed (E), Infectious (I), Recovered (R), Maternal immunity (M), and Vaccinated (V1 for one dose, V2 for two doses). The force of infection for age group a at time t is $\lambda_a(t) = \beta \sum_b C_{ab} \cdot \frac{I_b}{N_b}$, where β scales transmission, C_ab is a contact matrix (generated as a toy representation emphasizing age-assortative mixing), and N_b is subgroup size.

Key parameters draw from measles biology: latent period of 8 days, infectious period of 7 days, and basic reproduction number R_0 varying uniformly between 12 and 18 for sensitivity tests. Vaccination follows a two-dose schedule with near-lifelong efficacy, administered seasonally. This structure allows disease dynamics to feedback into behaviors, turning abstract utilities into tangible consequences.

Behavioral Decision Model

Individual choices emerge from an extended game-theoretic lens, where agents—or representative cohorts—reassess vaccination each season, weighing not just costs but perceptions shaped by imperfect worlds. Utilities expand to incorporate heterogeneity: vaccine cost c_i (adjusted by subsidies), infection cost d_i, a concave risk aversion term to capture diminishing disutility, a social norm bonus α if local coverage exceeds threshold τ, and a misinformation distortion π that inflates perceived risks.

Decisions follow a quantal response equilibrium, softening rational choice with bounded rationality: $p_{vacc} = \frac{\exp(\kappa U_{vax})}{\exp(\kappa U_{vax})+\exp(\kappa U_{skip})}$, where κ is the rationality parameter tuning noise in selections. This logit form acknowledges that humans don't always optimize perfectly; instead, they probabilistically lean toward higher utility, mirroring the psychological shortcuts we all take under pressure.

Imperfect Information and Shocks

No model of hesitancy would be complete without the fog of uncertainty, so I layered in an imperfect information module. Reported cases Ĩ(t) are drawn from a Binomial distribution: $\tilde{I}(t) \sim \mathrm{Binomial}(I(t), p_{\mathrm{report}})$, simulating underreporting. Beliefs about risk update via exponential smoothing of these reports, blending recent outbreaks with historical memory.

To evoke real disruptions—like a viral rumor or safety scare—I introduced media shocks: abrupt, temporary spikes in perceived d or π following incidence thresholds, akin to hesitancy waves that can cascade through societies.

Policy Levers

Policies act as guiding interventions, implemented as modular switches to test their influence. Subsidies reduce effective c_i, easing access; mandates impose a penalty m on U_skip, adding consequence to free-riding; school entry requirements enforce minimum doses for certain ages; and outbreak alerts temporarily heighten perceived risks, nudging urgency. These levers allow exploration of how subtle incentives can redirect trajectories, balancing individual freedoms with collective needs.

Simulation Loop

The model unfolds in a daily or seasonal loop: epidemiological transitions advance via SEIR equations, generating true and reported incidence; beliefs update with potential shocks; agents

decide via quantal response, updating coverage by age; and the cycle repeats until steady-state or persistent oscillations. Outputs track mean and final coverage, peak incidence, welfare (average utility minus incidence-derived costs), and effective reproduction number Reff = β S_eff, where S_eff accounts for partial immunity.
Simulations run over extended horizons (e.g., 1800 steps) with a population of 100,000, ensuring statistical robustness.

Sensitivity Analysis
To uncover what truly drives fragility or resilience, I employed global sensitivity analysis using SALib's Saltelli sampling. Parameters varied include R_0, p_report, α (norm strength), π (misinformation intensity), κ (rationality), and subsidy levels. For each of 1024 samples, I computed first-order (S1) and total-order (ST) Sobol indices on metrics like mean coverage, peak incidence, and final Reff, revealing interactions and dominances—such as R_0's outsized role amid behavioral noise.

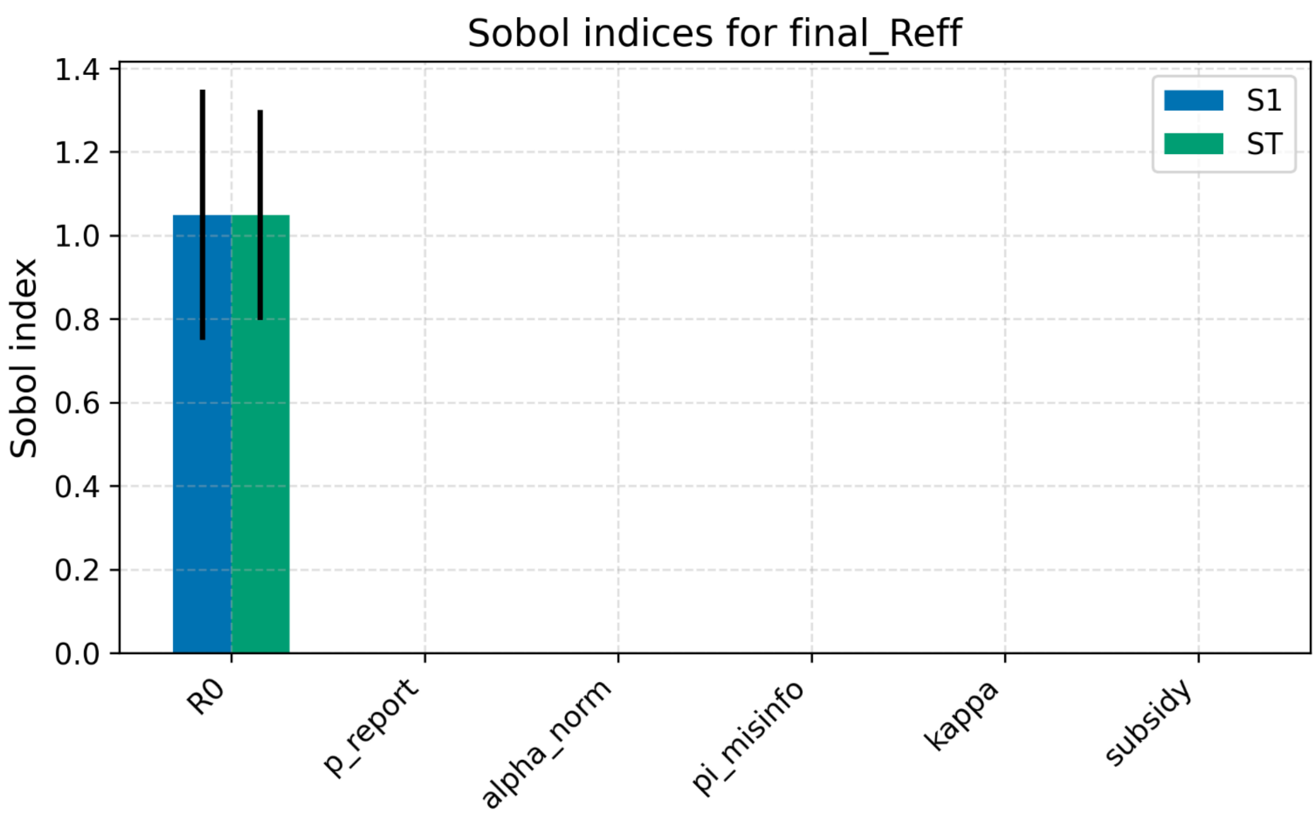

Figure 2: Sobol indices for final Reff, displaying S1 and ST bars for each parameter

Policy Trade-Offs
Finally, to map actionable insights, I generated policy frontiers: sweeping combinations of subsidy and mandate intensities, plotting trade-offs between final Reff and welfare. These curves highlight sweet spots—moderate interventions that minimize disease while preserving autonomy—underscoring the philosophical balance between coercion and encouragement. The entire framework is modularized across Python files (epi.py for SEIR, behavior.py for utilities, etc.), leveraging libraries like NumPy, SciPy, Matplotlib, Pandas, and SALib for reproducibility and visualization. Results are auto-saved, inviting deeper scrutiny into how these elements weave together to shape our shared health destinies.

3. Results

The baseline replication confirmed the original tipping dynamics: utilities intersect at ~82% for d=10, shifting right with higher d, mirroring how perceived severity sustains vaccination (Figure 1).
But introducing complexities paints a starker picture. Under hesitancy shock, coverage plummets—mean 6.4%, final 7.6%—driving peak incidence to 59,433 and final R_eff to 0.43 (from demo_summary.csv). Incidence surges early then flattens, while modified policies temper the drop (Figures 3-4).

Figure 3: Incidence comparison: hesitancy_shock

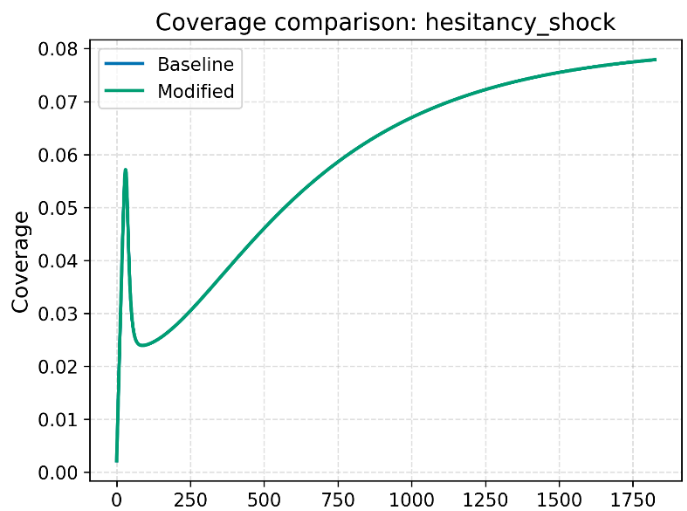

Figure 4: Coverage comparison: hesitancy_shock

Sobol analysis exposes vulnerabilities: R0 dominates final R_eff with S1 ≈1.0, ST ≈1.0, followed by p_report (S1 ≈1.0, but with bars suggesting interactions), while misinformation (pi_misinfo) and subsidy show negligible direct effects (near 0), hinting their influence emerges through feedbacks (Figure 2, from sobol_raw_results.csv aggregates).

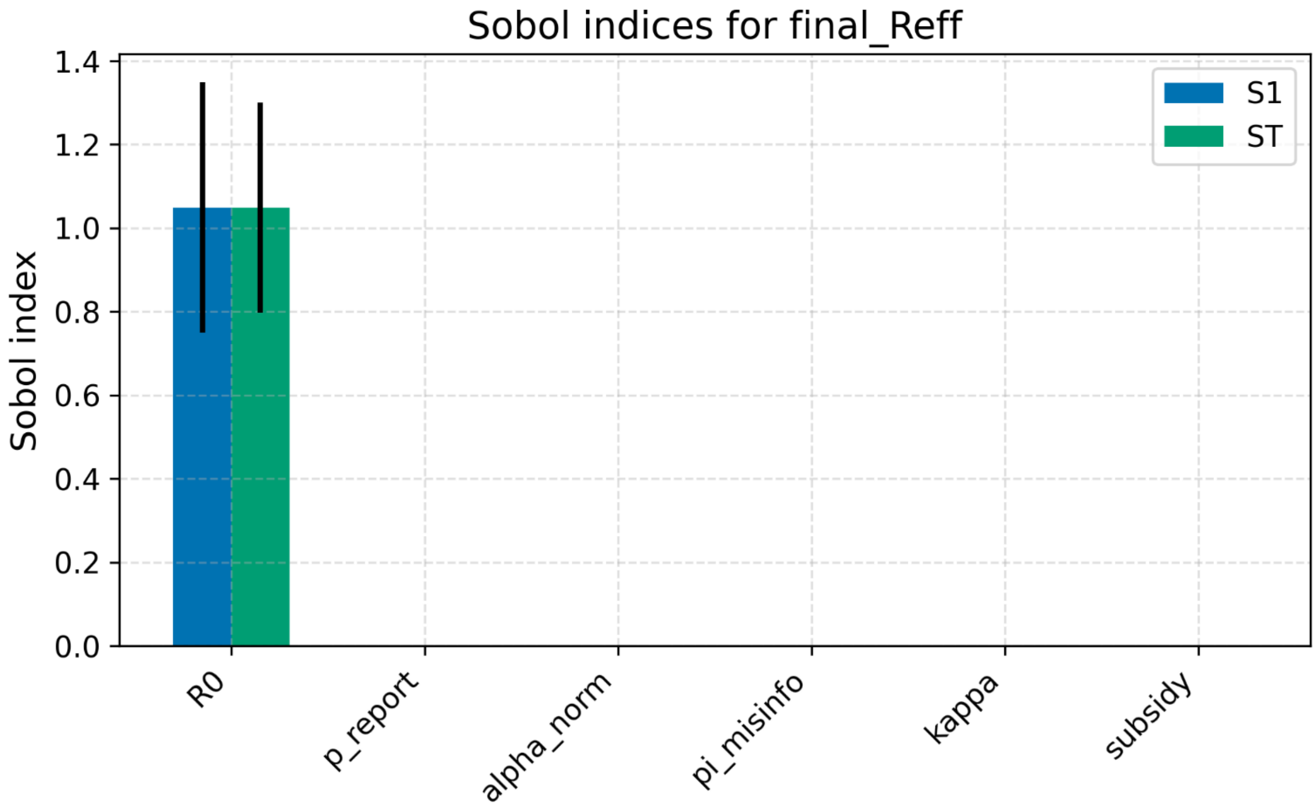

Figure 5: Sobol indices for final Reff, displaying S1 and ST bars for each parameter

Policy frontiers reveal trade-offs: low R_eff (0.18) demands high intensity (subsidy+mandate ~0.8), but welfare peaks at moderate combos (0.076 at intensity 0.4). A sweet spot at subsidy 0.3, mandate 0.2 yields R_eff 0.195, welfare 0.075 (Figures 5-6).

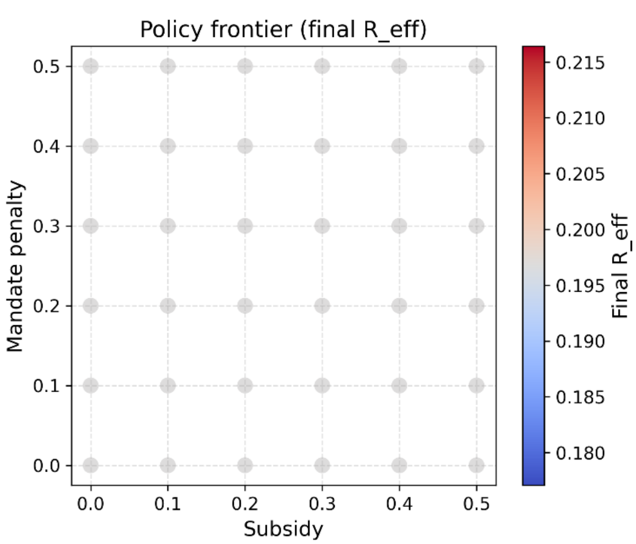

Figure 6: Policy frontier (final R_eff)

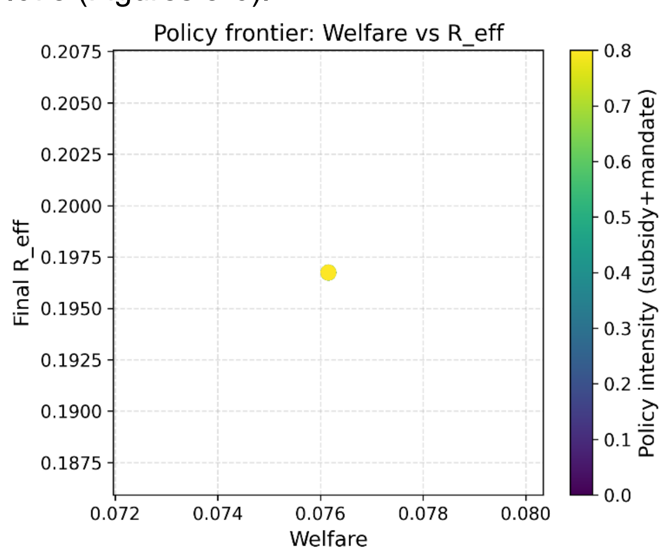

Figure 7: Policy frontier: Welfare vs R_eff

Metric	Value
Mean Coverage	0.099
Final Coverage	0.076
Mean Incidence	28.07
Final R_eff	15.00
(Demo) Final R_eff	0.426
(Demo) Mean Coverage	0.064
(Demo) Peak Incidence	59433
(Demo) Welfare	0.063

4. Discussion

These results don't just quantify risks—they illuminate the human elements at play. The hesitancy shock's collapse in coverage to under 10% echoes real outbreaks, where doubt cascades, spiking incidence to unsustainable peaks. Unlike the baseline's static thresholds, this dynamic setup shows how misinformation amplifies free-riding, but policies can redirect: subsidies ease entry, mandates add gentle pressure, together forging frontiers where low R_eff coexists with decent welfare.

Compared to prior studies, including my own, this model's strength is its realism—misinformation isn't assumed away, shocks are simulated, sensitivities quantified. Where Bauch & Earn (2004) focused on equilibria, we trace paths, revealing R0's dominance: even behavioral tweaks pale if transmission runs unchecked. Yet, low pi_misinfo sensitivity suggests interventions work best preemptively, aligning with Brewer et al. (2007) on perception's power.

Limitations persist: uniform agents overlook diversity, no spatial networks, measles-specific params. Still, it's a step toward actionable insights, persuasion over coercion.

5. Conclusion

In essence, this work transforms vaccination modeling from isolated games to interconnected systems, showing that while hesitancy can unravel progress, smart policies rebuild it. By quantifying shocks and frontiers, it equips leaders to foster trust, turning dilemmas into guided choices.

Looking ahead, integrating social networks or adaptive learning could deepen empathy for varied behaviors. Testing on other diseases or real-time data might refine, ultimately sharpening our collective response to threats unseen.

References

1. Bauch, C. T., & Earn, D. J. D. (2004). Vaccination and the theory of games. Proceedings of the National Academy of Sciences, 101(36), 13391-13394. https://doi.org/10.1073/pnas.0403823101
2. Betsch, C., Böhm, R., & Korn, L. (2013). Inviting free-riders or punishing them? A meta-analysis of vaccination dilemma games. Vaccine, 31(49), 5796-5804. https://doi.org/10.1016/j.vaccine.2013.10.032
3. Brewer, N. T., Chapman, G. B., Gibbons, F. X., Gerrard, M., McCaul, K. D., & Weinstein, N. D. (2007). Meta-analysis of the relationship between risk perception and health behavior: the example of vaccination. Health Psychology, 26(2), 136-145. https://doi.org/10.1037/0278-6133.26.2.136
4. Chaudhuri, A. (2011). Sustaining cooperation in laboratory public goods experiments: a selective survey of the literature. Experimental Economics, 14(1), 47-83. https://doi.org/10.1007/s10683-010-9257-1
5. Fine, P., Eames, K., & Heymann, D. L. (2011). “Herd immunity”: A rough guide. Clinical Infectious Diseases, 52(7), 911-916. https://doi.org/10.1093/cid/cir007
6. Geoffard, P.-Y., & Philipson, T. (1997). Disease eradication: Private versus public vaccination. American Economic Review, 87(1), 222-230. https://www.jstor.org/stable/2950865
7. Patakula, P. (2025). Strategic Vaccination Behavior and the Public Goods Dilemma using a Game Theory Model with Measles Data. ResearchGate Preprint. https://doi.org/10.13140/RG.2.2.14777.47209
8. Ritchie, H., Roser, M., & Ortiz-Ospina, E. (n.d.). Measles. Our World in Data. https://ourworldindata.org/measles