

Smooth or Polarized Behavior? A Mean-Field Infection Game^{*}

Audrey Hu[†]

City University of Hong Kong

Yating Yuan[‡]

University of Warwick

^{*}This paper supersedes the theoretical part of Hu, Han and Liu (2022), which was circulated under the working title "Polarized Public Responses to a Pandemic: Theory and Evidence". We are much indebted to Jim Bergin, Piotr Dworczak, Jacob Glazer, Steve Matthews, Motty Perry, Andy Skrzypacz, Wenji Xu, and Liang Zou for their invaluable comments, inputs, and advice on an earlier version of this paper. Our gratitude also extends to Daron Acemoglu, Andy Atkeson, Te Bao, Jan Boone, Dongkyu Chang, Yi-Chun Chen, Francesc Dilme, Christian Gollier, Ilwoo Hwang, Shenzhe Jiang, Yang Jiao, Jinwoo Kim, Nenad Kos, Takashi Kunimoto, Stephan Lauermann, Jiangtao Li, Jingfeng Lu, Yue Ma, Theo Offerman, Masao Ogaki, Marco Ottaviani, Alessandro Pavan, Sven Rady, Philip Reny, Yangbo Song, Dezső Szalay, Jingyi Xue, Zaifu Yang, Andy Zapechelnnyuk, Wenbo Zhao, Yu Zheng, Junjie Zhou, the seminars participants of Chinese University of Hong Kong, National University of Singapore, Osaka University, Seoul National University, Singapore Management University, University of Bonn and York University, and conference participants at Asian meeting of the Econometrics society (Tokyo) and Conference on Mechanism and Institution Design (Singapore), The 11th Conference of Hong Kong Micro Theory Group (2023), the RES Annual Conference (2024), and the Conference on Mechanism and Institution Design (2024) for their helpful comments and feedback. The usual disclaimer applies.

[†]9-256, Lau Ming Wai Academic Building, Department of Economics and Finance, City University of Hong Kong, Hong Kong SAR. Telephone number: +852 34426767. Email: audrey.hu@cityu.edu.hk.

[‡]Department of Economics, University of Warwick, Coventry CV47AL, United Kingdom. Email: yating.yuan@warwick.ac.uk

Abstract

We study a discrete-time mean-field game of epidemic behavior with heterogeneous, atomless agents who choose communicable activity in response to an evolving infection state. The aggregate state augments prevalence with the cross-section of remaining susceptibles, allowing for consistent population dynamics. We prove existence of a pure Markov mean-field equilibrium (MFE) and characterize equilibrium behavior via a simple curvature test that compares the log-slopes of activity benefits and infection risk. When marginal utility declines at least as fast as marginal risk, equilibrium choices are smooth and obey first-order KKT conditions. When marginal risk decays faster (a “front-loaded risk” environment), any interior stationary point is a local minimum and equilibrium choices are bang–bang: a unique cutoff partitions types into zero-activity and full-activity camps. A random-matching Poisson micro-foundation delivers a clean rate comparison—smooth vs polarization—and implies an endogenous transition from polarized to smooth responses as prevalence falls. The framework rationalizes heterogeneous public reactions across outbreaks and yields operational diagnostics for targeted, state-contingent policy.

Keywords: mean-field equilibrium; epidemic behavior; polarization; bang–bang control; random matching; Poisson transmission; heterogeneous agents; curvature test.

JEL: C73; C61; D83; I18.

1 Introduction

Why do public health crises sometimes elicit sharply divided behaviors—with some people taking extreme precautions while others behave as usual—whereas in other episodes public responses look comparatively smooth? Survey evidence from the early COVID-19 period suggests pronounced polarization in preventive choices across many cities and countries (e.g., YouGov), while similar divides were far less prominent in other outbreaks.¹ A natural conjecture is that polarized behavior merely mirrors polarized preferences or demographics. This paper shows that polarization can emerge *endogenously* even when the distribution of types is atomless and unimodal, preferences are smooth, and beliefs are common: the transmission technology alone can generate sharply divided equilibrium responses.

We develop a dynamic, discrete-time mean-field model of epidemic behavior with heterogeneous agents who differ only in *vulnerability* (the lifetime utility if infected). At the start of each period, a susceptible agent observes the aggregate infectious state and chooses a level of communicable activity. Infection risk depends on both the infectious state and own activity. If infected, the agent becomes immune thereafter. In this environment we study *pure Markov mean-field equilibria* (MFE): each agent best responds to a conjectured aggregate law of motion, and in equilibrium the population’s behavior reproduces that same law.

Our contributions are threefold.

1. *MFE formulation and existence with an augmented state.* We cast the epidemic interaction as a mean-field game with the *augmented aggregate state* (ϑ, m) : the fraction infectious ϑ and the cross-section of remaining susceptibles m (a sub-probability measure over types). We allow relaxed (mixed) Markov policies and prove the existence of a *pure* Markov MFE by (i) solving each agent’s Bellman equation for any continuous aggregate transition, (ii) showing continuity of the best-response map, and (iii) applying Schauder/Kakutani on a compact, convex policy space. This delivers a fixed point in policies whose induced law of motion coincides with the conjectured one. Uniqueness is not required for our behavioral characterization.
2. *Two equilibrium regimes: smooth vs. polarized.* We provide a crisp curvature condition that

¹For example, during the 2009 H1N1 outbreak, surveys in Saudi Arabia, Australia, and India reported distributions of precautionary behavior that were not bimodal; see (Balkhy et al., 2010; Seale et al., 2009; Kamate et al., 2010).

separates two qualitative regimes of equilibrium behavior, stated in terms of the log-slope (curvature over slope) of benefits and risks:

$$\frac{v''(a)}{v'(a)} \begin{matrix} \leq \\ \geq \end{matrix} \frac{p_{22}(\vartheta, a)}{p_2(\vartheta, a)} \quad \text{for all } a \in [0, 1].$$

When the marginal utility from activity declines at least as fast as the marginal infection risk (left inequality), optimal actions are *smooth*: the KKT first-order condition pins down a unique interior choice except when boundaries bind. When the marginal infection risk decays *faster* with action than marginal utility (right inequality), any interior stationary point is a *local minimum*; equilibrium choices are therefore *bang-bang* (polarized). We show existence of a unique cutoff $c(\vartheta, m)$ such that higher-vulnerability agents (lower x) optimally choose 0 while lower-vulnerability agents choose 1.

3. *Micro-foundation via random matching and Poisson transmission.* We micro-found the infection technology with random matching and Poisson arrival conditional on exposure. With $p(\vartheta, a) = \vartheta(1 - e^{-\lambda a})$ and an exponentially smoothed activity utility $v(a) = (1 - e^{-\gamma a})/(1 - e^{-\gamma})$, the regime test collapses to a comparison of *rates*:

$$\text{smooth if } \gamma \geq \lambda \quad \text{and} \quad \text{polarized if } \gamma < \lambda.$$

In the latter (“front-loaded risk”) case, the unique interior MB=MC crossing is a local minimum, so optimal choices jump to corners and a cutoff policy obtains. As prevalence falls along the equilibrium path, the economy can transition endogenously from a polarized to a smooth regime.

The key insight is that discrete-time decisions interact with the curvature of the transmission technology. When the instantaneous hazard of infection decays quickly with additional activity (relative to the decay in marginal utility), a myopic interior crossing of marginal benefit and marginal cost *cannot* be optimal: the objective is locally convex at that point. The global maximizers are at the extremes, producing a cutoff in types and a polarized aggregate response. Conversely, when utility’s marginal benefit decays at least as fast as marginal risk, the objective is locally concave and the interior FOC characterizes behavior.

Relation to the literature. Our work contributes to mean-field models of strategic behavior and to the economic-epidemiology literature that endogenizes preventive activity (Toxvaerd, 2020; Farboodi et al., 2021; Carnehl et al., 2023; McAdams et al., 2023; Egorov et al., 2021; Keppo et al., 2021; Baril-Tremblay et al., 2021; Dasaratha, 2023). While Acemoglu et al. (2023) show that extreme actions may arise from network formation, we obtain sharp polarization in a representative, frictionless matching environment driven by the memoryless (Poisson) nature of exposure risk. On the methodological side, our augmented-state MFE existence result follows the fixed-point approach standard in mean-field games (continuous time and discrete time), adapted here to a discrete-time SIR setting with a measure-valued state.

Section 2 sets up the model and defines MFE with the augmented state. Section 3.2 proves existence of a pure Markov MFE. Section 4 characterizes smooth and polarized regimes via the curvature test and establishes the cutoff structure in the polarized case. Section 5 provides the random-matching Poisson micro-foundation and two extensions (endogenous transmission intensity and resource-based prevention). Section 6 concludes.

2 Model

Consider a society populated by a continuum of agents with a population size equal to 1. Time is divided into discrete periods $t \in \mathbb{N} = \{1, 2, \dots\}$ of equal length. An infectious disease occurs at $t = 0$ and evolves over time. Using the SIR terminology, at the start of any period t (or epoch t), agents can be partitioned into three subpopulations: the *susceptible*, the *infectious*, and the *recovered*. The susceptible agents have not been previously infected. The infectious agents were infected in the previous period $t - 1$ and spread the disease in period t ; they will become uninfected and immune to the disease from epoch $t + 1$ onward. Let the size of the infectious population be denote by $\vartheta_t \in \Theta := [0, 1]$, referred to as the *infectious state* at epoch t . Our focus will be on the behavior of the susceptible agents.

Endogenized infection probability. At each epoch $t \in \mathbb{N}$, a susceptible agent chooses a level of *communicable activity* (or *action*) $a_t \in A := [0, 1]$ for period t . The agent will then be infected with probability $p(\vartheta_t, a_t)$ by the end of period t . We assume that the function $p : \Theta \times X \rightarrow [0, 1]$ is twice continuously differentiable: $p \in C^2(\Theta \times X)$. Let the subscripts

denote the partial derivatives with respect to the associated arguments. We assume

$$p(\vartheta, 0) = p(0, a) = p_1(\vartheta, 0) = p_2(0, a) \quad \forall (\vartheta, a), \text{ and } p(\vartheta, 1) < 1 \quad \forall \vartheta > 0, \quad (1)$$

and that $p_1(\vartheta, a) > 0 \quad \forall a > 0$, $p_2(\vartheta, a) > 0 \quad \forall \vartheta > 0$, and $p_{12}(\vartheta, a) \geq 0$. Thus, increasing either ϑ_t or a_t increases the infection risk, and the marginal probability of infection w.r.t. action is also (weakly) increasing in the infectious state. At the end of the period, agents observe whether they have become infected.

Heterogeneous agents. Each agent has a type $x \in [0, 1]$. The type space $X := [0, 1]$ is endowed with its Borel σ -algebra. Types are distributed according to a (bounded) Borel probability measure $m_0 \in \mathcal{P}(X)$, where $\mathcal{P}(X)$ denotes the set of Borel probability measures on X that are bounded from ∞ . We do *not* require m_0 to admit a density; m_0 may be discrete, continuous, or mixed.

With $\delta \in (0, 1)$ being the common discount factor, each agent enjoys a lifetime utility $U = 1/(1 - \delta)$ from full activity without infection. If a type- x agent becomes infected, they will suffer an expected disutility $D(x)$,² reducing U to $B(x) = U - D(x)$. Specifically, we assume $B : X \rightarrow (0, U]$ is continuously differentiable: $B \in C^1(X)$ and strictly increasing: $B'(x) > 0$ for all $x \in X$. Thus, lower types expect higher losses upon infection, e.g., elderly people.

Payoffs. Let $\mathbf{a} = \{a_t\}_{t=1}^\infty$ represent an activity plan by the type- x susceptible agents, in which a_t is a measurable function of all information at epoch t (dropping the dependence of a_t on x to simplify notation). The flow utility from engaging in communicable activities is measured by a twice continuously differentiable function $v : A \rightarrow [0, 1]$, satisfying $v' > 0$ and $0 \leq v(0) < v(1) = 1$. Being infected in period $t \in \mathbb{N}$ means that they will realize a discounted payoff (viewed at epoch 1):

$$\sum_{s=1}^t \delta^{s-1} v(a_s) + \delta^t B(x)$$

Let $\boldsymbol{\vartheta} = \{\vartheta_t\}_{t=1}^\infty$ represent an aggregate infection path. Given $(\boldsymbol{\vartheta}, \mathbf{a})$, the probability of not being infected through $t-1$ periods equals $\prod_{s=1}^{t-1} (1 - p(\vartheta_s, a_s))$ (defining $\prod_{s=1}^0 (\cdot) \equiv 1$). Therefore,

²The disutility can be related to any consequences of being sick of the disease, ranging from mild discomfort to death. The simplifying assumption here is that each type x of the agents is able to rationally quantify their expected disutility beforehand, conditional on being infected.

their discounted expected payoff equals

$$\sum_{t=1}^{\infty} \delta^{t-1} \prod_{s=1}^{t-1} (1 - p(\vartheta_s, a_s)) [v(a_t) + \delta p(\vartheta_t, a_t) B(x)] \quad (2)$$

The type- x agents' objective is to choose an activity plan that maximizes the expected payoff in (2), as long as they remain uninfected.

3 Population Dynamics

Because infections remove susceptible mass over time in a type-dependent way, it is convenient to work directly with sub-probability measures rather than $\mathcal{P}(X)$. Let $\Sigma(X)$ denote the Borel σ -algebra over X , and define

$$M := \{m : \Sigma(X) \rightarrow \mathbb{R}_+ \text{ such that } m \leq m_0\},$$

i.e., the set of sub-probability measures on X , endowed with the weak topology.

- **Aggregate state.** At time t , the aggregate state is the pair $(\vartheta_t, m_t) \in \Theta \times M$, where:
 - ϑ_t is the current fraction of infected (or the infectious state);
 - m_t is the cross-sectional distribution of *susceptible* types at time t (a sub-probability measure with total mass $m_t(X) \leq 1$).
- **Initial condition.** In period 0, the aggregate state $(\vartheta_0, m_0) \in [0, 1] \times \mathcal{P}(X)$ is assumed to be given, where m_0 is the initial cross-section of susceptible types. Since agents were unaware of the disease, they took action 1 in period 0 as usual. By the end of period 0, the population observes ϑ_1 of the agents infected, given by

$$\vartheta_1 = \int_X p(\vartheta_0, 1) m_0(dx)$$

As a result, the mass of type- x agents who remain susceptible is reduced to:

$$m_1(dx) = (1 - p(\vartheta_0, 1)) m_0(dx).$$

- **Population update under stationary Markov policies.** A (possibly randomized) stationary Markov policy is a stochastic kernel

$$\mu(\cdot \mid x, \vartheta, m) \in \Delta(A),$$

measurable in (x, ϑ, m) . A *pure* policy $\alpha : X \times \Theta \times M \rightarrow A$ is the special case where

$$\mu(du \mid x, \vartheta, m) = \delta_{\alpha(x, \vartheta, m)}(du),$$

with δ_a denoting the Dirac mass at a .

Given a stationary Markov policy μ and current aggregate state (ϑ_t, m_t) , next period's aggregate state is

$$\Phi_\mu(\vartheta_t, m_t) := \vartheta_{t+1} = \int_X \int_A p(\vartheta_t, a) \mu(da \mid x, \vartheta_t, m_t) m_t(dx),$$

and the law of motion for m_t is given by

$$\langle \varphi, \Psi_\mu(\vartheta_t, m_t) \rangle := \int_X \varphi(x) \left(1 - \int_A p(\vartheta_t, a) \mu(da \mid x, \vartheta_t, m_t) \right) m_t(dx), \quad \forall \varphi \in C(X),$$

which defines $(\Phi_\mu, \Psi_\mu) : \Theta \times M \rightarrow \Theta \times M$.

Thus, under μ the mass of type- x agents who remain susceptible is reduced by the infection factor

$$1 - \int_A p(\vartheta_t, a) \mu(da \mid x, \vartheta_t, m_t).$$

A pure policy α is recovered by taking $\mu(\cdot \mid x, \vartheta, m)$ degenerate at $\alpha(x, \vartheta, m)$, in which case (??)–(??) are special cases of the formulas above.

3.1 Individual dynamic program.

Fix a continuous aggregate transition $S = (\Phi, \Psi) : \Theta \times M \rightarrow \Theta \times M$. The value function of a type- x agent, when the aggregate evolves according to S , is a function

$$V_S : X \times \Theta \times M \rightarrow \mathbb{R}$$

that solves the Bellman equation

$$V_S(x, \vartheta, m) = \max_{a \in A} \left\{ v(a) + \delta p(\vartheta, a) B(x) \right. \quad (3)$$

$$\left. + \delta(1 - p(\vartheta, a)) V_S(x, \Phi(\vartheta, m), \Psi(\vartheta, m)) \right\}. \quad (4)$$

Because $\delta \in (0, 1)$ and $0 \leq 1 - p(\vartheta, a) \leq 1$, the Bellman operator associated with (3) is a contraction in the sup norm, and thus admits a unique bounded fixed point V_S .

Theorem 1 (Individual problem: existence under a given aggregate law) *For each continuous aggregate transition $S = (\Phi, \Psi)$, there exists a unique bounded solution V_S to the Bellman equation (3). Moreover, there exists a stationary Markov policy (possibly mixed) $\mu_S(\cdot \mid x, \vartheta, m)$ attaining this value, in the sense that for every (x, ϑ, m) the support of $\mu_S(\cdot \mid x, \vartheta, m)$ is contained in the set of maximizers of the right-hand side of (3).*

Proof. Fix a continuous aggregate transition $S = (\Phi, \Psi) : \Theta \times M \rightarrow \Theta \times M$ and define the Bellman operator T_S on the space of bounded functions

$$\mathcal{B} := \{V : X \times \Theta \times M \rightarrow \mathbb{R} \text{ bounded and Borel measurable}\}$$

by

$$(T_S V)(x, \vartheta, m) := \max_{a \in A} \left\{ v(a) + \delta p(\vartheta, a) B(x) + \delta(1 - p(\vartheta, a)) V(x, \Phi(\vartheta, m), \Psi(\vartheta, m)) \right\}.$$

Step 1: T_S is a contraction. Equip $\mathcal{B}$ with the sup norm

$$\|V\|_\infty := \sup_{(x, \vartheta, m) \in X \times \Theta \times M} |V(x, \vartheta, m)|.$$

Let $V_1, V_2 \in \mathcal{B}$. For any (x, ϑ, m) and any $a \in A$,

$$\begin{aligned} & \left| \left[v(a) + \delta p(\vartheta, a) B(x) + \delta(1 - p(\vartheta, a)) V_1(x, \Phi(\vartheta, m), \Psi(\vartheta, m)) \right] \right. \\ & \quad \left. - \left[v(a) + \delta p(\vartheta, a) B(x) + \delta(1 - p(\vartheta, a)) V_2(x, \Phi(\vartheta, m), \Psi(\vartheta, m)) \right] \right| \\ & \leq \delta |1 - p(\vartheta, a)| |V_1(x, \Phi(\vartheta, m), \Psi(\vartheta, m)) - V_2(x, \Phi(\vartheta, m), \Psi(\vartheta, m))| \\ & \leq \delta \|V_1 - V_2\|_\infty, \end{aligned}$$

using $0 \leq 1 - p(\vartheta, a) \leq 1$. Taking the maximum over $a \in A$ on both sides yields

$$|(T_S V_1)(x, \vartheta, m) - (T_S V_2)(x, \vartheta, m)| \leq \delta \|V_1 - V_2\|_\infty$$

for all (x, ϑ, m) . Hence

$$\|T_S V_1 - T_S V_2\|_\infty \leq \delta \|V_1 - V_2\|_\infty,$$

so T_S is a contraction on $(\mathcal{B}, \|\cdot\|_\infty)$ with modulus $\delta \in (0, 1)$. By Banach's fixed point theorem, there exists a unique bounded function $V_S \in \mathcal{B}$ such that

$$V_S = T_S V_S,$$

i.e. V_S is the unique bounded solution of the Bellman equation (3).

Step 2: existence of an optimal stationary Markov policy. Fix $(x, \vartheta, m) \in X \times \Theta \times M$ and consider the one-step objective as a function of a :

$$F_{x, \vartheta, m}(a) := v(a) + \delta p(\vartheta, a) B(x) + \delta (1 - p(\vartheta, a)) V_S(x, \Phi(\vartheta, m), \Psi(\vartheta, m)), \quad a \in A.$$

By assumption, v and p are continuous in a , B is continuous in x , and V_S is bounded. Moreover, Φ and Ψ are continuous in (ϑ, m) by assumption on S . Therefore, for fixed (x, ϑ, m) , the function $a \mapsto F_{x, \vartheta, m}(a)$ is continuous on the compact action set $A = [0, 1]$.

By the Weierstrass theorem, $F_{x, \vartheta, m}$ attains its maximum on A . Define the argmax set

$$\Gamma_S(x, \vartheta, m) := \arg \max_{a \in A} F_{x, \vartheta, m}(a),$$

which is nonempty and compact for each (x, ϑ, m) .

Define a relaxed stationary Markov policy μ_S by specifying that, for each (x, ϑ, m) , the probability measure $\mu_S(\cdot \mid x, \vartheta, m) \in \Delta(A)$ has support contained in $\Gamma_S(x, \vartheta, m)$. For instance, one may choose

$$\mu_S(\cdot \mid x, \vartheta, m) := \delta_{a^*(x, \vartheta, m)}(\cdot),$$

where $a^*(x, \vartheta, m) \in \Gamma_S(x, \vartheta, m)$ is any Borel-measurable selection.³ By construction, μ_S is a

³The existence of such a measurable selection follows from standard measurable selection theorems (e.g. Kuratowski–Ryll–Nardzewski) applied to the argmax correspondence Γ_S .

stationary Markov policy, and at every state (x, ϑ, m) it places its mass on maximizers of $F_{x, \vartheta, m}$. Therefore, μ_S attains the value V_S . ■

3.2 Mean-Field Equilibrium

In our setting there is a continuum of infinitesimal agents, each facing the same dynamic infection environment but having no individual impact on it. A mean-field equilibrium (MFE) formalizes this idea: each agent takes as given an aggregate law of motion for the epidemic (how the population infection state (ϑ_t, m_t) evolves over time), chooses an optimal *Markov policy* $\mu^*(\cdot \mid x, \vartheta_t, m_t)$ given that law, and in equilibrium the population's actual behavior under μ^* reproduces exactly the same law of motion. In other words, in an MFE beliefs about the aggregate dynamics are correct and each agent optimally responds to those dynamics.

This equilibrium concept is particularly appropriate for our epidemic model for three reasons. First, agents are atomless, so it is natural for them to treat the aggregate state as exogenous when making individual decisions. Second, infection and recovery dynamics are driven by the distribution of behaviors across types rather than by any single agent, which is exactly what the "mean field" captures. Third, MFE allows us to focus on tractable Markov policies that depend only on the current aggregate state (ϑ_t, m_t) while still enforcing full consistency between individual optimization and population-level epidemic dynamics.

Definition 1 (Markov Mean-Field Equilibrium with augmented state) *A pair (V^*, μ^*) is a (relaxed) Markov mean-field equilibrium (MFE) if:*

1. Optimality. *Let $S^* := (\Phi_{\mu^*}, \Psi_{\mu^*})$ be the aggregate transition induced by μ^* . Then V^* is the unique bounded solution to the Bellman equation*

$$V^*(x, \vartheta, m) = \max_{a \in [0, 1]} \left\{ v(a) + \delta p(\vartheta, a) B(x) + \delta (1 - p(\vartheta, a)) V^*(x, S^*(\vartheta, m)) \right\},$$

and for all (x, ϑ, m) the relaxed policy $\mu^(\cdot \mid x, \vartheta, m)$ is supported on the set of maximizers of the right-hand side, i.e.*

$$\text{supp } \mu^*(\cdot \mid x, \vartheta, m) \subseteq \Gamma_{S^*}(x, \vartheta, m),$$

where $\Gamma_{S^}(x, \vartheta, m)$ denotes the argmax set.*

2. Consistency. For every initial aggregate state (ϑ_0, m_0) , the induced path $(\vartheta_t, m_t)_{t \geq 0}$ satisfies

$$\vartheta_{t+1} = \Phi_{\mu^*}(\vartheta_t, m_t) = \int_X \int_A p(\vartheta_t, a) \mu^*(da \mid x, \vartheta_t, m_t) m_t(dx),$$

and

$$\langle \varphi, m_{t+1} \rangle = \langle \varphi, \Psi_{\mu^*}(\vartheta_t, m_t) \rangle = \int_X \varphi(x) \left(1 - \int_A p(\vartheta_t, a) \mu^*(da \mid x, \vartheta_t, m_t) \right) m_t(dx), \quad \forall \varphi \in C(X),$$

for all $t = 0, 1, 2, \dots$. Equivalently, $(\vartheta_{t+1}, m_{t+1}) = S^*(\vartheta_t, m_t)$ for all t .

3.3 Equilibrium existence

We first establish the existence of a MFE.

Theorem 2 (Existence of Markov MFE with augmented state) *Let $X = [0, 1]$ be the type space, and let $\Theta = [0, 1]$ be the space of aggregate infection levels. Let M be the set of Borel sub-probability measures on X , endowed with the weak topology. The aggregate state space is $S := \Theta \times M$, which is compact and metrizable.*

For any stationary Markov policy $\mu(\cdot \mid x, \vartheta, m)$, define the induced aggregate transition $S_\mu = (\Phi_\mu, \Psi_\mu) : S \rightarrow S$ as above. Then there exists a Markov mean-field equilibrium (V^, μ^*) with augmented state (ϑ, m) in the sense of Definition 1.*

Proof (sketch). Fix any state-transition kernel $S : S \rightarrow S$. By Theorem 1, the associated Bellman operator T_S is a contraction on the space of bounded functions $V : X \times S \rightarrow \mathbb{R}$ endowed with the sup norm, and there exists a unique bounded solution V_S to the Bellman equation. For each (x, ϑ, m) , let $\Gamma_S(x, \vartheta, m) \subset A$ denote the (nonempty, compact) argmax set in the Bellman problem with continuation S .

An individually optimal relaxed policy given S is any Markov kernel $\mu(\cdot \mid x, \vartheta, m)$ whose support at each state (x, ϑ, m) is contained in $\Gamma_S(x, \vartheta, m)$. By Berge's maximum theorem and the continuity of primitives, the argmax correspondence Γ_S is upper hemicontinuous with nonempty compact values, and the induced best-response correspondence from aggregate transitions S to optimal kernels μ is nonempty, convex-valued, and upper hemicontinuous on a compact convex set of Markov kernels.

Conversely, given a kernel μ , the induced aggregate transition $S_\mu = (\Phi_\mu, \Psi_\mu)$ is continuous in (ϑ, m) because it is defined by integrating continuous functions of (x, ϑ, m, a) against μ and m on compact domains. Thus, composing the maps “ $\mu \mapsto S_\mu$ ” and “ $S \mapsto \text{optimal } \mu$ ” yields a self-correspondence on the compact convex set of Markov kernels that satisfies the conditions of Kakutani’s fixed point theorem.

Therefore there exists a kernel μ^* such that μ^* is an optimal response to S_{μ^*} . Let $S^* := S_{\mu^*}$ and $V^* := V_{S^*}$. By construction, V^* solves the Bellman equation with continuation S^* , and μ^* is pointwise optimal given V^* , establishing optimality. The definition of S_μ ensures that the aggregate path (ϑ_t, m_t) induced by μ^* satisfies $(\vartheta_{t+1}, m_{t+1}) = S^*(\vartheta_t, m_t)$ for all t , which is exactly the consistency condition. Hence (V^*, μ^*) is a Markov MFE with augmented state (ϑ, m) . ■

4 Equilibrium Behavior

In this section we characterize the mean-field equilibrium (MFE) under two qualitatively different parameter regimes. Although our existence result is formulated for (possibly randomized) relaxed stationary Markov policies, we will focus here on *pure* stationary Markov policies

$$\alpha^* : X \times \Theta \times M \rightarrow [0, 1], \quad (x, \vartheta, m) \mapsto \alpha^*(x, \vartheta, m),$$

which arise when the equilibrium relaxed policy assigns probability one to a single action at each state:

$$\mu^*(du \mid x, \vartheta, m) = \delta_{\alpha^*(x, \vartheta, m)}(du).$$

Let $S^* = (\Phi_{\mu^*}, \Psi_{\mu^*})$ denote the aggregate transition induced by the equilibrium policy, and write $S^*(\vartheta, m) = (\vartheta', m')$. It is convenient to define, for a type- x susceptible agent,

$$D^*(x, S^*(\vartheta, m)) := V^*(x, S^*(\vartheta, m)) - B(x),$$

the *infection premium*: the difference between the ex-ante continuation value of being currently susceptible and the utility $B(x)$ the agent would obtain if already infected and immune. By the one-stage deviation principle, in an MFE the pure policy α^* must satisfy, for each (x, ϑ, m) , that

$\alpha^*(x, \vartheta, m)$ maximizes

$$F_{x, \vartheta, m}(a) := v(a) + \delta p(\vartheta, a)B(x) + \delta(1 - p(\vartheta, a))V^*(x, S^*(\vartheta, m)), \quad a \in [0, 1]. \quad (5)$$

4.1 Monotonicity in type

We first establish how the equilibrium value and infection premium vary with the health type x .

Lemma 1 *For each aggregate state (ϑ, m) , both $V^*(x, S^*(\vartheta, m))$ and $D^*(x, S^*(\vartheta, m))$ are differentiable in x , and*

$$\frac{\partial}{\partial x} V^*(x, S^*(\vartheta, m)) \geq 0, \quad \frac{\partial}{\partial x} D^*(x, S^*(\vartheta, m)) < 0, \quad \forall x \in X.$$

Proof. Fix a type x and consider the equilibrium path $(\vartheta_t, m_t)_{t \geq 0}$ under α^* , with $(\vartheta_t, m_t) = S^*(\vartheta_t, m_t)$ along this path. For brevity, denote

$$\alpha_t^* := \alpha^*(x, \vartheta_t, m_t), \quad p_t := p(\vartheta_t, \alpha_t^*).$$

From the Bellman equation under S^* , we have

$$V^*(x, \vartheta_t, m_t) = v(\alpha_t^*) + \delta p_t B(x) + \delta(1 - p_t) V^*(x, S^*(\vartheta_t, m_t)).$$

Subtracting $B(x)$ from both sides and using $D^*(x, S^*) = V^*(x, S^*) - B(x)$ yields the recursion

$$D^*(x, \vartheta_t, m_t) = v(\alpha_t^*) - (1 - \delta)B(x) + \delta(1 - p_t) D^*(x, S^*(\vartheta_t, m_t)).$$

Iterating this recursion and writing $\prod_{r=0}^{-1}(\cdot) \equiv 1$, we obtain the sequence representation

$$D^*(x, \vartheta_t, m_t) = \sum_{s=0}^{\infty} \delta^s \prod_{r=0}^{s-1} (1 - p_{t+r}) \left[v(\alpha_{t+r}^*) - (1 - \delta)B(x) \right].$$

Using the envelope theorem (or differentiating term-by-term and applying dominated convergence), we may differentiate with respect to x , treating the optimal controls α_t^* as fixed:

$$\frac{\partial}{\partial x} D^*(x, \vartheta_t, m_t) = -(1 - \delta) B'(x) \sum_{s=0}^{\infty} \delta^s \prod_{r=0}^{s-1} (1 - p_{t+r}).$$

Each factor satisfies $0 \leq 1 - p_{t+r} \leq 1$, so

$$0 < \sum_{s=0}^{\infty} \delta^s \prod_{r=0}^{s-1} (1 - p_{t+r}) \leq \sum_{s=0}^{\infty} \delta^s = \frac{1}{1 - \delta},$$

and hence

$$0 < (1 - \delta) \sum_{s=0}^{\infty} \delta^s \prod_{r=0}^{s-1} (1 - p_{t+r}) \leq 1.$$

Because $B'(x) > 0$ by assumption, it follows that

$$-B'(x) \leq \frac{\partial}{\partial x} D^*(x, S^*(\vartheta_t, m_t)) < 0,$$

so D^* is strictly decreasing in x and its derivative is bounded below by $-B'(x)$.

Finally, recall that

$$D^*(x, S^*(\vartheta_t, m_t)) = V^*(x, S^*(\vartheta_t, m_t)) - B(x),$$

so

$$\frac{\partial}{\partial x} D^*(x, S^*(\vartheta_t, m_t)) = \frac{\partial}{\partial x} V^*(x, S^*(\vartheta_t, m_t)) - B'(x).$$

Combining this with $-B'(x) \leq \partial_x D^* < 0$ implies

$$0 \leq \frac{\partial}{\partial x} V^*(x, S^*(\vartheta_t, m_t)) < B'(x),$$

so V^* is nondecreasing in x . This proves the lemma. ■

We next establish how the equilibrium strategies vary with x .

Lemma 2 *The equilibrium strategy $\alpha^*(x, \vartheta, m)$ is nondecreasing in x .*

Proof. Take any $x \in X$. The cross-partial derivative of $F_{x,\vartheta,m}(a)$, holding $\alpha_t^* = \alpha^*(x, \vartheta_t, m_t)$ fixed, is given by

$$\frac{\partial^2}{\partial a \partial x} F_{x,\vartheta,m}(a)|_{a=\alpha^*} = -\delta p_2(\vartheta, \alpha^*) \frac{\partial}{\partial x} D^*(x, S^*(\vartheta, m)) \geq 0 \quad (6)$$

where the inequality derives from Lemma 1. Standard arguments then establish the conclusion of the lemma (e.g., Milgrom and Shannon (1994)). ■

4.2 Smooth public response to the health crisis

We first analyze a smooth case in which agents trade off marginal benefits of activity against marginal infection risk in an interior way. The key condition is that the per-period objective in a is (weakly) concave at an interior optimum.

Proposition 1 (Smooth response) *Suppose that, for all $a \in [0, 1]$ and $\vartheta \in \Theta$,*

$$\frac{v''(a)}{v'(a)} \leq \frac{p_{22}(\vartheta, a)}{p_2(\vartheta, a)}. \quad (7)$$

(i) *The equilibrium policy α^* satisfies the KKT conditions:*

$$\begin{aligned} v'(\alpha^*) - \delta p_2(\vartheta, \alpha^*) D^*(x, S^*(\vartheta, m)) &= 0 \quad \text{if } \alpha^*(x, \vartheta, m) \in (0, 1) \\ v'(1) - \delta p_2(\vartheta, 1) D^*(x, S^*(\vartheta, m)) &\geq 0 \quad \text{iff } \alpha^*(x, \vartheta, m) = 1 \\ v'(0) - \delta p_2(\vartheta, 0) D^*(x, S^*(\vartheta, m)) &\leq 0 \quad \text{iff } \alpha^*(x, \vartheta, m) = 0 \end{aligned} \quad (8)$$

(ii) *Along any equilibrium path $(\vartheta_t, m_t)_{t \geq 0}$, for each $x \in X$,*

$$\alpha^*(x, \vartheta_t, m_t) \rightarrow 1 \quad \text{as } t \rightarrow \infty.$$

Proof. (i) Fix (ϑ, m) and a type x . Because individual agents are atomless, a deviation in a does not affect the aggregate transition S^* , so the continuation value $V^*(x, S^*(\vartheta, m))$ can be treated as constant in the choice of a . Differentiating (5) with respect to a gives

$$\begin{aligned} \frac{\partial}{\partial a} F_{x,\vartheta,m}(a) &= v'(a) + \delta p_2(\vartheta, a) B(x) - \delta p_2(\vartheta, a) V^*(x, S^*(\vartheta, m)) \\ &= v'(a) - \delta p_2(\vartheta, a) D^*(x, S^*(\vartheta, m)). \end{aligned}$$

Similarly,

$$\frac{\partial^2}{\partial a^2} F_{x,\vartheta,m}(a) = v''(a) - \delta p_{22}(\vartheta, a) D^*(x, S^*(\vartheta, m)).$$

At an interior optimum $a = \alpha^*(x, \vartheta, m) \in (0, 1)$ we have $F'_{x,\vartheta,m}(\alpha^*) = 0$, that is,

$$v'(\alpha^*) = \delta p_2(\vartheta, \alpha^*) D^*(x, S^*(\vartheta, m)).$$

Substituting into F'' yields

$$F''_{x,\vartheta,m}(\alpha^*) = v'(\alpha^*) \left(\frac{v''(\alpha^*)}{v'(\alpha^*)} - \frac{p_{22}(\vartheta, \alpha^*)}{p_2(\vartheta, \alpha^*)} \right).$$

Under condition (7), the term in parentheses is nonpositive, and since $v' > 0$, we obtain $F''_{x,\vartheta,m}(\alpha^*) \leq 0$, with strict inequality if (7) is strict at α^* . Thus interior optima satisfy the first line in (8) and are (weak) local maxima. The boundary conditions follow from the usual Kuhn-Tucker argument: if $\alpha^* = 1$ were not optimal, a small deviation inward would strictly increase $F_{x,\vartheta,m}$, contradicting optimality; the same reasoning applies at 0.

(ii) Along any equilibrium path, the mass of newly infected agents in period t is ϑ_t . Since each infected agent remains infectious for exactly one period and then becomes immune, the total mass of agents who are ever infected is $\sum_{t=1}^{\infty} \vartheta_t \leq 1$. Hence $\vartheta_t \rightarrow 0$ as $t \rightarrow \infty$.

By continuity of p_2 and the assumptions in (1) that $p_2(\vartheta, 1) \rightarrow 0$ as $\vartheta \downarrow 0$, we have $p_2(\vartheta_t, 1) \rightarrow 0$ as $t \rightarrow \infty$. The equilibrium value functions are uniformly bounded, so there exists $K < \infty$ with $|D^*(x, S^*(\vartheta_t, m_t))| \leq K$ for all t and x . Then

$$v'(1) - \delta p_2(\vartheta_t, 1) D^*(x, S^*(\vartheta_t, m_t)) \rightarrow v'(1) > 0 \quad \text{as } t \rightarrow \infty,$$

so by the boundary KKT condition in (8), for all sufficiently large t the unique maximizer satisfies $\alpha^*(x, \vartheta_t, m_t) = 1$. Hence $\alpha^*(x, \vartheta_t, m_t) \rightarrow 1$ as $t \rightarrow \infty$. ■

4.3 Polarized public response to the health crisis

We now turn to a regime in which the equilibrium response is *polarized* (bang-bang): agents either shut down activity completely or behave as if there were no epidemic. This occurs when the per-period objective in a is (strictly) convex rather than concave at stationary points.

Proposition 2 (Polarized response) *Suppose that, for all $a \in [0, 1]$ and $\vartheta \in \Theta$,*

$$\frac{v''(a)}{v'(a)} > \frac{p_{22}(\vartheta, a)}{p_2(\vartheta, a)}. \quad (9)$$

Then:

(i) *There exists a cutoff function $c : \Theta \times M \rightarrow [0, 1]$ such that the equilibrium policy is bang-bang:*

$$\alpha^*(x, \vartheta, m) = \begin{cases} 1 & \text{if } x \geq c(\vartheta, m) \\ 0 & \text{if } x < c(\vartheta, m) \end{cases}$$

(ii) *Along any equilibrium path (ϑ_t, m_t) , we have $c(\vartheta_t, m_t) \rightarrow 0$ as $t \rightarrow \infty$. In fact, for t large enough, $c(\vartheta_t, m_t) = 0$ so that $\alpha^*(x, \vartheta_t, m_t) \equiv 1$ for all x .*

Proof. (i) Under (9), at any interior point $a \in (0, 1)$ satisfying $F'_{x, \vartheta, m}(a) = 0$, we have

$$F''_{x, \vartheta, m}(a) = v'(a) \left(\frac{v''(a)}{v'(a)} - \frac{p_{22}(\vartheta, a)}{p_2(\vartheta, a)} \right) > 0,$$

so a is a *local minimum* rather than a maximum of $F_{x, \vartheta, m}$. Hence no interior point can be optimal, and for any (x, ϑ, m) we must have

$$\alpha^*(x, \vartheta, m) \in \{0, 1\}.$$

To determine which boundary is chosen, compute the difference in value between the two corners:

$$\begin{aligned} \Delta(x, \vartheta, m) &:= F_{x, \vartheta, m}(1) - F_{x, \vartheta, m}(0) \\ &= \left[v(1) + \delta p(\vartheta, 1)B(x) + \delta(1 - p(\vartheta, 1))V^*(x, S^*(\vartheta, m)) \right] \\ &\quad - \left[v(0) + \delta p(\vartheta, 0)B(x) + \delta(1 - p(\vartheta, 0))V^*(x, S^*(\vartheta, m)) \right]. \end{aligned}$$

Using $p(\vartheta, 0) = 0$ and simplifying,

$$\Delta(x, \vartheta, m) = v(1) - v(0) - \delta p(\vartheta, 1)D^*(x, S^*(\vartheta, m)).$$

By Lemma 1, $D^*(x, S^*(\vartheta, m))$ is strictly decreasing in x , so $\Delta(x, \vartheta, m)$ is strictly increasing in x . Define the cutoff $c(\vartheta, m)$ as follows:

- If there exists $x \in X$ with $\Delta(x, \vartheta, m) = 0$, let $c(\vartheta, m)$ be that unique value; - If $\Delta(x, \vartheta, m) > 0$ for all x , set $c(\vartheta, m) := 0$ (then all types choose 1); - If $\Delta(x, \vartheta, m) < 0$ for all x , set $c(\vartheta, m) := 1$ (then all types choose 0).

In all cases, the optimal policy is

$$\alpha^*(x, \vartheta, m) = \begin{cases} 1 & \text{if } x \geq c(\vartheta, m) \\ 0 & \text{if } x < c(\vartheta, m) \end{cases}$$

as claimed.

(ii) Along any equilibrium path (ϑ_t, m_t) , we have $\vartheta_t \rightarrow 0$ as $t \rightarrow \infty$, by the same argument as in Proposition 1: the total mass of agents who are ever infected is at most 1, and each infected agent is infectious for only one period. Assume $p(\vartheta, 1) \rightarrow 0$ as $\vartheta \downarrow 0$ (which holds, for example, if $p(\vartheta, a) = \vartheta q(a)$ for some q).

For each (ϑ_t, m_t) , the cutoff $c(\vartheta_t, m_t)$ is defined (in the interior case) by the equation

$$\Delta(c(\vartheta_t, m_t), \vartheta_t, m_t) = v(1) - v(0) - \delta p(\vartheta_t, 1) D^*(c(\vartheta_t, m_t), S^*(\vartheta_t, m_t)) = 0.$$

The infection premium $D^*(x, S^*(\vartheta_t, m_t))$ is uniformly bounded in t and x because V^* is bounded and $B(x)$ lies in $(0, U]$. Thus there exists $K < \infty$ with

$$|D^*(x, S^*(\vartheta_t, m_t))| \leq K, \quad \forall x, t.$$

Since $p(\vartheta_t, 1) \rightarrow 0$, we have

$$\frac{v(1) - v(0)}{\delta p(\vartheta_t, 1)} \rightarrow +\infty \quad \text{as } t \rightarrow \infty.$$

But

$$\Delta(x, \vartheta_t, m_t) = v(1) - v(0) - \delta p(\vartheta_t, 1) D^*(x, S^*(\vartheta_t, m_t)) \geq v(1) - v(0) - \delta p(\vartheta_t, 1) K,$$




Figure 1: In a strong infectious environment, agents' reactions are polarized (blue curves). In a mild infectious environment, agents' reactions are smooth.

so for t large enough, $\Delta(x, \vartheta_t, m_t) > 0$ for all x . In that case $c(\vartheta_t, m_t) = 0$, i.e. all types strictly prefer $\alpha^* = 1$. Consequently $c(\vartheta_t, m_t) \rightarrow 0$ as $t \rightarrow \infty$, and for all sufficiently large t the equilibrium policy satisfies

$$\alpha^*(x, \vartheta_t, m_t) \equiv 1 \quad \text{for all } x \in X.$$

■

Figure 1 illustrates the different public reactions to the two regimes.

5 Micro Foundation of Equilibrium Behavior

This section provides an intuitive foundation for Propositions 1 (smooth response) and 2 (polarized response). We model infection as the outcome of random matching with the currently infectious fraction and Poisson transmission over the time spent in communicable activities.

5.1 Random matching and Poisson transmission

At the beginning of period t , each susceptible agent chooses an activity duration $a \in [0, 1]$ (e.g., time spent in social settings). Let $\vartheta_t \in [0, 1]$ be the fraction infectious at the start of period t . Under random matching, a susceptible is *exposed* to an infectious environment with probability

ϑ_t . Conditional on exposure, infection arrives according to a Poisson process with constant rate $\lambda > 0$ over the interval of length a . Thus the conditional infection probability $q(a)$ solves

$$q'(a) = \lambda(1 - q(a)), \quad q(0) = 0,$$

with solution $q(a) = 1 - e^{-\lambda a}$. The *unconditional* infection probability before time a is therefore

$$p(\vartheta_t, a) = \vartheta_t(1 - e^{-\lambda a}).$$

(For more generality one may replace ϑ_t by an increasing exposure map $g(\vartheta_t)$; the exposition below is unchanged.)

For utility from activity, we use a parallel exponential form that nests diminishing marginal utility and delivers clean comparative statics. Let $v \in C^2([0, 1])$ satisfy

$$v'(a) = \gamma(C - v(a)), \quad v(0) = 0,$$

so $v(a) = C(1 - e^{-\gamma a})$. Choosing $C = (1 - e^{-\gamma})^{-1}$ normalizes $v(1) = 1$, hence

$$v(a) = \frac{1 - e^{-\gamma a}}{1 - e^{-\gamma}}, \quad v'(a) = \frac{\gamma e^{-\gamma a}}{1 - e^{-\gamma}}.$$

Both $p(\vartheta, a)$ and $v(a)$ are strictly increasing and strictly concave in a .

Local MB–MC tradeoff. Fix an aggregate state (ϑ_t, m_t) and a type x . Under the mean-field policy, the agent's infection premium is

$$D^*(x, S^*(\vartheta_t, m_t)) := V^*(x, S^*(\vartheta_t, m_t)) - B(x),$$

the continuation value difference between being currently susceptible and being already infected/immune. A myopic marginal extension of activity from a to $a + da$ yields:

Marginal benefit (MB): $v'(a) da$.

Marginal cost (MC): the extra infection probability $p_2(\vartheta_t, a) da = \vartheta_t \lambda e^{-\lambda a} da$ times the con-




Figure 2: Strong vs. mild infectious regime.

tinuation loss $\delta D^*(x, S^*)$,

$$MC(a; x) \approx \delta \vartheta_t \lambda e^{-\lambda a} D^*(x, S^*) da.$$

Thus the first-order condition (FOC) at an interior optimum $\alpha^*(x, \vartheta_t, m_t) \in (0, 1)$ matches Proposition 1:

$$v'(\alpha^*) = \delta p_2(\vartheta_t, \alpha^*) D^*(x, S^*).$$

Two infectious regimes. The shapes of MB and MC over a govern equilibrium behavior. For the exponential forms above,

$$\frac{v''(a)}{v'(a)} = -\gamma, \quad \frac{p_{22}(\vartheta, a)}{p_2(\vartheta, a)} = -\lambda.$$

Hence the curvature comparison in Propositions 1 and 2 collapses to a comparison of the rates γ and λ (see Figure 2):

- **Mild infectious regime** ($\lambda \leq \gamma$). MB (v') declines weakly *faster* than MC ($\delta p_2 D^*$) as a grows; the per-period objective is (weakly) concave at interior points, and the KKT system admits an interior solution when the boundary conditions do not bind. This is exactly the “smooth” case: agents pick a unique interior $\alpha^*(x, \vartheta, m) \in (0, 1)$ determined by $MB = MC$, unless pushed to a boundary.

- **Strong infectious regime** ($\lambda > \gamma$). In this environment MC declines *faster* with a than MB does. Any stationary point that solves $MB = MC$ is a *local minimum* (since $F'' > 0$ at that point). Thus the maximizer must be at a boundary: $\alpha^* \in \{0, 1\}$. Comparing the two corners yields

$$F_{x,\vartheta,m}(1) - F_{x,\vartheta,m}(0) = (v(1) - v(0)) - \delta p(\vartheta, 1) D^*(x, S^*),$$

which is strictly increasing in x because D^* is strictly decreasing in x (Lemma 1). This produces a unique cutoff $c(\vartheta, m)$ with bang-bang behavior:

$$\alpha^*(x, \vartheta, m) = \begin{cases} 1 & \text{if } x \geq c(\vartheta, m) \\ 0 & \text{if } x < c(\vartheta, m) \end{cases}$$

Why polarization is a “natural” response in discrete time? In discrete time, a small marginal step in a need not reveal the global shape of the objective. When $\lambda > \gamma$, starting from $a = 0$ the MC can exceed MB locally until the crossing point (a “danger zone”, see the shaded area on Figure 2(right) for type- x' agents), yet the *global* payoff at $a = 1$ can be strictly higher than at $a = 0$ because MC decays more steeply than MB as a increases. This is why the optimal choice jumps to a corner rather than settling at a local crossing point. The exponential case makes this stark: $v''/v' = -\gamma$, $p_{22}/p_2 = -\lambda$, so $\lambda > \gamma$ guarantees any interior crossing is a local minimum.

5.2 Two useful variations

Endogenous transmission intensity. Suppose the Poisson rate depends on prevalence, $\lambda = \lambda(\vartheta)$, strictly increasing in ϑ . Then

$$p(\vartheta, a) = 1 - e^{-\lambda(\vartheta)a}, \quad \frac{p_{22}}{p_2} = -\lambda(\vartheta).$$

For a given period t , the period- t regime (smooth vs. polarized) is governed by

$$\text{smooth if } \gamma \geq \lambda(\vartheta_t), \quad \text{polarized if } \gamma < \lambda(\vartheta_t).$$

Because $\vartheta_t \downarrow 0$ along any equilibrium path, a strongly infectious phase (polarized behavior) can endogenously transition to a mild phase (smooth interior behavior) as the epidemic wanes. If v is linear, the cutoff logic persists each period regardless of ϑ_t .

Resource-based prevention decisions. Let agents devote a share $u_t \in [0, 1]$ of a per-period resource to prevention (e.g., masking, filtration), reducing the instantaneous infection rate from λ to $(1 - u_t)\lambda$ over the period. The unconditional infection probability becomes

$$p(\vartheta_t, u_t) = \vartheta_t \left(1 - e^{-\lambda(1-u_t)}\right).$$

Let the period resource cost be $c(u_t) := 1 - v(1 - u_t)$ (so higher prevention reduces activity utility). The per-period objective is

$$1 - c(u_t) + \delta p(\vartheta_t, u_t) B(x) + \delta (1 - p(\vartheta_t, u_t)) V^*(x, S^*).$$

The FOC at an interior optimum is

$$c'(u_t) = \delta \vartheta_t \lambda e^{-\lambda(1-u_t)} D^*(x, S^*),$$

and the same curvature comparison applies with the map $u \mapsto 1 - u$:

$$\frac{d^2}{du^2} \log c'(u) \begin{matrix} \leq \\ \geq \end{matrix} -\lambda \quad \Longleftrightarrow \quad \text{smooth (interior) vs. polarized (corner) prevention.}$$

Equivalently, with $c'(u) = v'(1 - u)$ and $c''(u) = -v''(1 - u)$, the period is smooth when

$$\frac{c''(u)}{c'(u)} \leq \lambda \quad (\text{all } u),$$

and polarized when the inequality reverses. The qualitative economic message is unchanged: the side (benefit or risk) with the steeper log-slope dominates the global shape of the period objective and determines whether behavior is interior or bang-bang.

6 Conclusion

This paper develops a discrete-time mean-field model of epidemic behavior in which heterogeneous, atomless agents optimally choose communicable activity in response to an evolving infection state. We formulate the interaction as a mean-field game with an *augmented state* (ϑ, m) —the current infection prevalence and the cross-sectional measure of remaining susceptibles—and prove the existence of a *pure Markov* mean-field equilibrium (MFE). The existence result is robust to general type distributions (no density is required) and accommodates relaxed policies while delivering pure best responses on equilibrium paths.

A central contribution is a sharp characterization of equilibrium behavior via a curvature test that compares the log-slopes of benefits and risks:

$$\frac{v''(a)}{v'(a)} \begin{matrix} \leq \\ > \end{matrix} \frac{p_{22}(\vartheta, a)}{p_2(\vartheta, a)} \quad \text{for } a \in [0, 1].$$

When the marginal utility from activity declines *at least as fast* as the marginal infection risk (left inequality), equilibrium choices are *smooth* and satisfy the usual first-order (KKT) conditions. When the marginal infection risk decays *faster* with additional activity than marginal utility (right inequality), any interior stationary point is a local minimum and equilibrium choices are *bang–bang*. In this polarized regime there exists a unique cutoff $c(\vartheta, m)$ such that more vulnerable types optimally choose zero activity while less vulnerable types choose one. Along equilibrium paths, as prevalence falls, the cutoff declines and the economy can transition endogenously from polarized to smooth behavior.

We provide a transparent micro-foundation for these regimes using random matching and Poisson transmission. With $p(\vartheta, a) = \vartheta(1 - e^{-\lambda a})$ and $v(a) = (1 - e^{-\gamma a})/(1 - e^{-\gamma})$, the regime test collapses to a comparison of *rates*: smooth behavior when $\gamma \geq \lambda$ and polarization when $\gamma < \lambda$. Intuitively, a large, memoryless hazard makes infection risk *front-loaded*: once an agent crosses a personal tolerance, additional activity adds relatively little risk but still yields utility—pushing optimal choices to the corners and generating a cutoff in types without requiring any ex ante polarization in preferences or demographics.

Our results rationalize why some epidemics exhibit sharp divides while others do not: the same society, with the same unimodal distribution of vulnerabilities, can display polarized behavior if the transmission hazard is sufficiently high relative to the curvature of activity benefits,

and smooth behavior otherwise. The augmented-state MFE highlights that polarization is a *dynamic* property: targeted data on prevalence and transmissibility (which influence p_2 and p_{22}) can forecast shifts between regimes and thus inform timing of interventions.

The model formalizes the case for *targeted* measures over uniform mandates. In polarized regimes, the equilibrium already sorts agents by vulnerability; policy can *support* efficient sorting (e.g., subsidizing low-cost protection and remote access for high-risk types, facilitating safe activity for low-risk types) rather than suppressing it with blunt restrictions. As prevalence recedes, the theory predicts a natural unwinding of extreme behaviors, suggesting that time-varying, state-contingent guidance may achieve public-health goals with lower economic and social costs. More broadly, monitoring the relative curvature *in situ*—via estimates of effective contact intensity (hazard) and observed marginal utilities of activity—provides an operational criterion for when to expect polarization and when to expect smooth adjustment.

We abstracted from learning, network structure, capacity constraints in healthcare, and strategic externalities beyond infection (e.g., congestion in testing or treatment). Extending the augmented-state MFE to incorporate belief updating, endogenous social links, or vaccination would allow us to study how information and policy co-evolve with behavior. Finally, a full welfare analysis—quantifying wedges between decentralized MFE and socially optimal allocations and characterizing implementable instruments—is an important next step.

Overall, by embedding epidemic decisions in a mean-field framework and isolating the curvature-based mechanism behind smooth versus polarized responses, the paper offers a unified explanation for heterogeneous public behavior across outbreaks and a tractable set of diagnostics for policy design.

References

- Acemoglu, D., Makhdoumi, A., Malekian, A., and Ozdaglar, A. (2023). Testing, voluntary social distancing, and the spread of an infection. *Operations Research*.
- Balkhy, H. H., Abolfotouh, M. A., Al-Hathloul, R. H., and Al-Jumah, M. A. (2010). Awareness, attitudes, and practices related to the swine influenza pandemic among the saudi public. *BMC infectious diseases*, 10:1–7.

- Baril-Tremblay, D., Marlats, C., and Ménager, L. (2021). Self-isolation. *Journal of Mathematical Economics*, 93:102483.
- Carnehl, C., Fukuda, S., and Kos, N. (2023). Epidemics with behavior. *Journal of Economic Theory*, 207:105590.
- Dasaratha, K. (2023). Virus dynamics with behavioral responses. *Journal of Economic Theory*, 214:105739.
- Egorov, G., Enikolopov, R., Makarin, A., and Petrova, M. (2021). Divided we stay home: Social distancing and ethnic diversity. *Journal of Public Economics*, 194:104328.
- Farboodi, M., Jarosch, G., and Shimer, R. (2021). Internal and external effects of social distancing in a pandemic. *Journal of Economic Theory*, 196:105293.
- Kamate, S. K., Agrawal, A., Chaudhary, H., Singh, K., Mishra, P., and Asawa, K. (2010). Public knowledge, attitude and behavioural changes in an indian population during the influenza a (h1n1) outbreak. *The Journal of Infection in Developing Countries*, 4(01):007–014.
- Keppo, J., Kudlyak, M., Quercioli, E., Smith, L., and Wilson, A. (2021). The behavioral si* model, with applications to the swine flu and covid-19 pandemics. *University of Wisconsin-Madison*.
- McAdams, D., Song, Y., and Zou, D. (2023). Equilibrium social activity during an epidemic. *Journal of Economic Theory*, 207:105591.
- Milgrom, P. and Shannon, C. (1994). Monotone comparative statics. *Econometrica*, pages 157–180.
- Seale, H., McLaws, M.-L., Heywood, A. E., Ward, K. F., Lowbridge, C. P., Van, D., Galton, J., and MacIntyre, C. R. (2009). The community’s attitude towards swine flu and pandemic influenza. *Medical Journal of Australia*, 191(5):267–269.
- Toxvaerd, F. M. (2020). Equilibrium social distancing. *Working Paper*.