Mathematics

AIMS Mathematics, 11(11): 36835–36869.
DOI: 10.3934/math.20261424
Received: 12 September 2026
Revised: 09 October 2026
Accepted: 12 October 2026
Published: 06 November 2026

https://www.aimspress.com/journal/Math

Research article

Evolutionary dynamics and stability of connected autonomous driving policies with information quality feedback

Shih-Lin Lin*

Graduate Institute of Vehicle Engineering, National Changhua University of Education, Changhua City, Taiwan

* **Correspondence:** Email: lin040@cc.ncue.edu.tw.

Abstract: Adaptive driving policies can increase information demand as their adoption changes interactions between traffic roles. This study developed a deterministic mean-field evolutionary model of adaptive and fixed policies in left-turning and straight-going connected autonomous vehicle populations, coupled to an endogenous information-quality state. Market penetration is an external parameter, distinct from the policy shares. The payoff coefficients represent aggregate behavioral and resource sensitivities rather than a calibrated intersection game. We established positive invariance, characterized feasible interior equilibria, and derived a relative-entropy sufficient condition for global asymptotic stability. A Jacobian-based local criterion identified stable equilibria outside that sufficient region. Signed cross-role interactions and a common monotone nonlinear information benefit were examined to clarify the scope of the stability argument. Boundary linearization gave local invasion criteria and candidate policy invasion thresholds. A synthetic parameter family linked small invasion rates to prolonged relaxation and algebraic decay at the policy invasion threshold. Numerical experiments examined these mechanisms across parameter variations, including cases beyond the global sufficient condition. Information recovery can change equilibrium quality, whereas uniform scaling of the quality equation preserves equilibria while altering transients. These results describe policy–information feedback in the specified model; traffic calibration and independent safety evaluation are needed for engineering interpretation.

Keywords: evolutionary game; connected autonomous vehicles; information quality; population dynamics; Lyapunov stability
Mathematics Subject Classification: 34D23, 37N40, 91A22

Abbreviations	
AV	autonomous vehicle
CAV	connected and automated vehicle
HV	human-driven vehicle
SVO	social value orientation
VANET	vehicular ad hoc network
ODE	ordinary differential equation
PDE	partial differential equation

1. Introduction

1.1. Background and motivation

Connected and automated vehicles (CAVs) interact with human drivers, whose preferences vary. At unprotected left-turn intersections, left-turning and straight-going vehicles have different passage requirements, motivating role-specific policy payoffs. Li et al. [1] used social value orientation (SVO) in a two-population evolutionary game to examine the implications of adaptive and fixed social preferences. Their traffic-interaction setting motivates the role distinction considered here.

Policy performance also depends on information conditions. Lin and Chen [2] studied non-line-of-sight radar perception at intersections; Lin and Wan [3] discussed the integration of intelligent connected vehicle technologies; and Lin [4] incorporated communication quality and edge latency into cooperative driving assessment. An adaptive policy may improve coordination payoffs while increasing information exchange and computational demand. Motivated by these interactions, this study introduces dynamic feedback between policy adoption and information quality to examine how demand-induced depletion and quality recovery affect equilibria.

1.2. Research questions and scope

This study considers the relative shares of adaptive and fixed policies within left-turning and straight-going CAV populations. Table 1 summarizes the notation used throughout the analysis. The model is deterministic, well-mixed, and delay-free. SVO motivates the policy classes, which are represented through aggregate payoffs; online preference estimation and trajectory planning are outside the model.

The exogenous parameter p affects the payoff coefficients, whereas x, y, and q evolve. The dimensionless state q summarizes information conditions relevant to policy payoffs. Following the strategy–environment perspective of Weitz et al. [5] and Tilman et al. [6], and distinguishing it from reward feedback with delay [7], the model examines interior coexistence, cross-role interactions, information-resource feedback, and policy invasion.

Table 1. Variables, parameters, and derived quantities.

Symbol	Definition	Role in the analysis
p	CAV market penetration	Exogenous scenario parameter
x	Adaptive policy share within the left-turning CAV population	Endogenous evolutionary state
y	Adaptive policy share within the straight-going CAV population	Endogenous evolutionary state
q	Effective information quality	Endogenous state with recovery and depletion
$q_{\max}$	Quality upper bound without adaptive-policy load	$0 < q_{\max} \leq 1$
a_i	Baseline adaptive-minus-fixed payoff difference	Real coefficient
b_i	Cross-role marginal payoff effect	Signed complementarity or competition
$c_i,\ d_i$	Quality benefit and own-role adoption-cost sensitivities	Positive coefficients
$r_i,\ \gamma$	Quality depletion and recovery coefficients	Positive rates
κ_i	Policy adjustment rate	Positive; affects transients
ω	Multiplier of the entire quality equation	Positive response-speed factor
$\mu_i,\ \nu_i$	Conventional fixed-payoff intercept and cross effect	Equation (3.1)
g_i	Adaptive-minus-fixed mean payoff	Equation (3.3)
D	Determinant of the linear policy-equilibrium equations	Equation (3.7)
$u_i,\ v_i$	Intercept and quality slope of equilibrium policy shares	Equation (3.7)
$A,\ B$	Coefficients of the scalar quality-equilibrium equation	Equation (3.7)
w_i	Weights matching information benefits to depletion	$w_i = r_i/c_i$
$S,\ \Delta$	Policy dissipation matrix and four times its determinant	Equations (3.10), (3.13)
G^*	Total quality relaxation coefficient at equilibrium	$\gamma + r_L x^* + r_S y^*$
$X,\ Y,\ Z$	Deviations from the interior equilibrium	Section 3.4.1
$\mathcal{B},\ V$	Bernoulli relative entropy and Lyapunov function	Equation (3.11)
$\sigma_i,\ H_j$	Interior replicator pre-factors and cubic coefficients	Equation (3.14); $j = 1,2,3$
$h,\ K,\ n$	Common quality-benefit function, shape scale, and Hill exponent	Proposition 3.2
$V_h,\ V_\omega$	Nonlinear-benefit and time-scaled Lyapunov functions	Equations (3.18), (3.25)
$I_i,\ \lambda_i$	Boundary invasion payoff and linear growth rate	$\lambda_i = \kappa_i I_i$
$p_{c,i}$	Candidate role-specific policy invasion threshold	Zero of $I_i(p)$
$\alpha_i,\ \beta_i$	Slope and intercept of the affine invasion indicator	Equation (3.23) only
$e_k,\ \varepsilon$	Sampled full-state error and tolerance	Equation (4.1)
$T_\varepsilon,\ T$	Sampled convergence time and observation horizon	Equation (4.1)
$\delta,\ \tau$	Distance from the synthetic threshold and spectral relaxation scale	Section 4.5.3

Coupling a strategy evolution to another changing system's property is an established coevolutionary-game approach. Perc and Szolnoki [8] reviewed such feedback through interaction networks, mobility, reputation, and other evolving attributes. Perc et al. [9] showed how structured interactions and collective dynamics shape cooperation, while Jusup et al. [10] connected traffic, cooperation, social networks, and intelligent machines within social physics. This broader literature motivates examining policy adoption together with its information environment. The present contribution is the equilibrium and stability analysis of a specific two-role recovery–depletion model; it does not introduce coevolution itself or resolve the spatial and stochastic mechanisms studied in those reviews.

Game-theoretic resource coordination also appears in energy systems. Cheng et al. [11] reviewed urban-energy models for demand response, load management, and adaptive decision making, while Wang et al. [12] reviewed evolutionary-game approaches to energy-storage coordination. Their common concern is how incentives and resource constraints shape interacting agents' decisions. These reviews provide cross-domain context; the present contribution is an explicit stability analysis of two policy populations coupled to an information-quality recovery–depletion state.

1.3. Objectives and main contributions

We propose a nonlinear ordinary differential equation model with three state variables. Its principal contributions are as follows:

1) A three-state mean-field model links role-specific policy payoffs to information recovery and adoption-induced depletion, with an explicit distinction between policy shares and market penetration.

2) Positive invariance, equilibrium feasibility, global sufficient stability, and local Jacobian stability are analyzed together, distinguishing a Lyapunov certificate from the larger region in which an equilibrium may remain locally stable.

3) Signed cross-role interactions and common monotone nonlinear information benefits clarify which features are essential to the Lyapunov argument. Uniform information-timescale scaling is treated as a corollary of that construction.

4) Synthetic numerical experiments examine parameter dependence and dynamics outside the global sufficient condition. Near a constructed policy invasion threshold, local analysis relates long transients to vanishing linear relaxation rates and nonlinear decay at criticality.

The model provides a theoretical account of policy adoption with information feedback. Its stability statements concern dynamical convergence and do not assert evolutionary stability or social-welfare optimality.

1.4. Organization of the paper

Section 2 reviews the related literature and positions the study. Section 3 develops the model and stability theorems, Section 4 presents numerical validation, and Section 5 summarizes the conclusions and directions for future research.

2. Related literature

2.1. Autonomous driving interactions and social value orientation

Schwarting et al. [13] incorporated SVO into behavioral inference and planning to represent how drivers weigh their own interests against those of others. A sociality probe, proposed by Liu et al. [14], uses game-theoretic inverse reinforcement learning to estimate individual and group payoff weights from interaction trajectories. These approaches provide interpretable representations and estimation methods for social preferences. Policy diffusion under information resource constraints, however, requires a separate evolutionary model.

Jusup et al. [10] placed traffic, cooperation, and human–machine interactions within a broader social-physics framework. This perspective motivates connecting policy adoption to its information environment, while the individual preference weights remain inputs to the aggregate utility representation.

Li et al. [1] connected traffic-role interactions, SVO-based payoffs, and evolutionary adoption. Their evolutionary shares describe AV/HV adoption within each traffic role. We retain the left-turn/straight-going role distinction and the idea that a policy's relative payoff depends on the interacting population. Here, however, x and y describe adaptive-policy shares within CAV populations, market penetration p is prescribed, and information quality q is an additional dynamic state. The present coefficients are not mapped one to one to their calibrated payoffs, and no reduction to their calibrated payoff system is established by fixing q. This is a conceptually motivated formulation rather than a formally nested extension; neither x nor y is an individual SVO angle.

2.2. Vehicular-network cooperation, perception, and edge computing

The editorial by Lin and Wan [3] provides background on the integration of perception, communication, and automation. Lin and Chen [2] studied around-the-corner non-line-of-sight perception at intersections using deep learning and frequency-modulated continuous-wave radar. Shivshankar and Jamalipour [15] applied evolutionary game theory to cooperation in vehicular ad hoc networks (VANETs) under different network conditions, linking those conditions to cooperation incentives.

Lin's [4] tail-risk allocation framework incorporates age of information, perception confidence, and edge latency into cooperative-driving assessment. These studies motivate a payoff dependence on information conditions. In the present formulation, q enters an aggregate payoff difference and is not itself a collision-risk or fairness measure.

2.3. Mixed-traffic stability and critical penetration

Ma et al. [16] analyzed string stability and the fundamental diagram of mixed traffic, while Chou et al. [17] examined the relationship between ring-road stability and the string stability of car-following models. Lin [18] incorporated anticipative headway adjustment and cooperative velocity-difference damping into an optimal-velocity model. Under the baseline parameters of that study, the critical penetration levels were approximately 0.173 for the long-wave approximation and 0.171 for the

finite ring system. Liard et al. [19] used a coupled PDE–ODE model to describe traffic flow and autonomous vehicle control.

These studies concern perturbations in velocity, headway, or density. The present analysis instead asks whether a rare policy can grow and whether policy shares and information quality return to equilibrium after perturbation. Its policy invasion thresholds arise from payoff differences near a boundary equilibrium and are conceptually distinct from traffic-wave stability thresholds.

2.4. Evolutionary games with environmental and resource feedback

The review by Perc and Szolnoki [8] places simultaneous evolution of strategies and other system properties within the established coevolutionary-game literature. Weitz et al. [5] and Tilman et al. [6] subsequently examined explicit strategy–environment feedback, while De Leenheer et al. [20] studied resource use and cooperation in a chemostat. In the present model, the coevolving property is shared information quality: adaptive-policy demand reduces it, recovery restores it, and its current value changes relative policy payoffs. The specific recovery–depletion law and matched benefit–load weights make the equilibrium characterization and Lyapunov cancellation analytically tractable.

Gong et al. [7] and Wang and Meng [21] analyzed cooperation dynamics with reward feedback and delay. Liu and Hu [22] studied logit dynamics and reverse-cyclic mutations in an asymmetric three-player, three-strategy game. The present model retains two policies per role to obtain verifiable stability conditions. Information response speed adjusts the current rate of change, whereas a delay introduces dependence on past states. The response-speed analysis therefore does not establish delay-induced Hopf bifurcations.

2.5. Spatial structure, finite populations, and adaptive interaction networks

Perc et al. [9] reviewed how spatial organization, group interactions, and self-organization influence cooperation. Ohtsuki and Nowak [23] studied replicator dynamics on graphs; Madeo et al. [24] analyzed self-loops and link removal; and Bratus et al. [25] examined explicit spatial structure and global regulation. Such formulations make interaction structure explicit, whereas the present mean-field model summarizes its effects through role-specific payoff and load coefficients.

Traulsen et al. [26] analyzed stochastic invasion and fixation in finite populations. Fontanari [27] showed that demographic noise can make deterministic coexistence metastable in collective-risk games. These findings motivate future comparisons between the present deterministic equilibria and finite-fleet policy dynamics.

Mohd and Park [28] studied environmental effects on cyclic competition, and Park and Jang [29] analyzed the structural stability of coexistence. Du et al. [30] considered reinforcement learning–driven adaptive networks, while Choi et al. [31] examined recognition error. These mechanisms suggest extensions beyond the role heterogeneity represented by the present coefficients.

2.6. Synthesis and positioning of the study

Table 2 summarizes the relevant approaches and the scope of the present study.

The proposed model combines two policy-adoption states with information recovery and depletion, allowing equilibrium feasibility, global sufficient stability, local stability, and boundary

invasion to be compared within one formulation. Habib [32] studied the use of transient dynamics near saddle-node bifurcations. Here, the relevant relaxation rates and critical decay are derived from the present equations, without assuming that its policy invasion threshold is a saddle-node bifurcation.

Table 2. Related research and the scope of the present study.

Research area	References	Existing basis	Treatment in this study
SVO and driving interactions	[1,10,13,14]	Preference inference and strategy evolution	Policy adoption within CAV populations
Mixed-traffic stability	[16–19]	Headway, velocity, and density perturbations	Distinction of policy invasion thresholds
Vehicular networks and cooperative perception	[2–4,15]	Perception, communication, and cooperation	Quality links payoffs and load
Environmental and resource feedback	[5–8,20]	Strategy–environment coevolution	Two-role recovery–depletion model
Delay and multiple strategies	[21,22]	Update rules and bifurcation	Theorems restricted to two policies
Spatial structure and finite populations	[9,23–27]	Network and stochastic effects	Limits on extrapolation from mean-field results
Heterogeneity and adaptive networks	[28–31]	Coexistence, recognition, and network adaptation	Role differences and future extensions

3. Theoretical framework and mathematical model

3.1. Conventional approaches and their limitations

In a two-population evolutionary game with fixed payoffs, policy shares change according to their payoff advantage over the population mean. For policy shares x, y, a standard form is

$$\dot{x} = x(1-x)(\mu_L + \nu_L y), \qquad \dot{y} = y(1-y)(\mu_S + \nu_S x). \tag{3.1}$$

In Eq (3.1), the constants μ_i, ν_i are determined by the fixed payoff matrices. This standard two-population structure represents evolution driven by relative payoffs. The traffic-role interpretation is motivated by Li et al. [1], while the additional information state introduced below makes policy payoffs depend on the recovery and depletion of information resources.

Evolutionary games with environmental feedback already couple strategies to an evolving environment [5,6], and reward feedback with delay has also been studied [7]. The present formulation combines two traffic-role policy populations with an aggregate information state in a structure that permits explicit equilibrium and stability analysis.

3.2. Coupled model of policy payoffs and information quality

3.2.1. States and payoff structure

Consider left-turning and straight-going roles, denoted by L and S. Each CAV population uses an adaptive policy A or a fixed policy F, with adaptive shares x and y, respectively. Effective information quality is q, with unloaded upper bound $q_{\max}$. Market penetration p is an exogenous scenario parameter. Adaptive means adjustment to the interacting population; it does not imply aggressive passage behavior.

With rows and columns ordered as adaptive and fixed policies, the role-specific relative payoff matrices are

$$\begin{aligned} \Pi_L(x,q) &= \begin{pmatrix} a_L + b_L + c_L q - d_L x & a_L + c_L q - d_L x \\ 0 & 0 \end{pmatrix}, \\ \Pi_S(y,q) &= \begin{pmatrix} a_S + b_S + c_S q - d_S y & a_S + c_S q - d_S y \\ 0 & 0 \end{pmatrix}. \end{aligned} \tag{3.2}$$

The rows of Eq (3.2) represent the focal policies and the columns the other role's policies. Subtracting the fixed-policy payoff separately in each column preserves the relative payoffs and replicator dynamics, so its zero entries are a normalization. The intercept a_i represents the adaptive-minus-fixed payoff difference at the zero-state reference. The coefficient b_i measures how adoption in the other role changes that difference: a positive value describes complementarity, while a negative value describes competition. The coefficient c_i measures the relative benefit of information quality, and d_i describes the marginal adoption-dependent cost within the focal role. Averaging over the other population gives

$$\begin{aligned} g_L(x,y,q) &= a_L + b_L y + c_L q - d_L x, \\ g_S(x,y,q) &= a_S + b_S x + c_S q - d_S y. \end{aligned} \tag{3.3}$$

Equation (3.3) is an affine aggregate payoff model. A possible calibration would first define a common utility scale using prespecified weights for travel time, comfort, risk, and social outcomes. Controlled experiments or microscopic simulations could then compare adaptive and fixed policies while varying the opponent-policy share, information quality, and focal adoption share. Regression of these utility differences would estimate a_i, b_i, c_i, and d_i; separate information-load measurements would inform the quality equation. Independent variation is necessary to distinguish correlated effects, and the utility scale and policy-update time scale must be fixed for identification. This is an estimation route, not a derivation from a calibrated SVO game; the present experiments use synthetic coefficients.

3.2.2. Evolutionary equations and parameter conditions

Policy shares follow two-strategy replicator dynamics. Information quality recovers toward $q_{\max}$ at rate γ and is depleted by demand associated with adaptive policies in both roles, yielding

$$\begin{aligned} \dot{x} &= \kappa_L x(1-x)(a_L + b_L y + c_L q - d_L x), \\ \dot{y} &= \kappa_S y(1-y)(a_S + b_S x + c_S q - d_S y), \\ \dot{q} &= \gamma(q_{\max} - q) - q(r_L x + r_S y). \end{aligned} \tag{3.4}$$

In Eq (3.4), κ_i controls policy adjustment, and r_i controls quality depletion. The direct costs $d_L x$ and $d_S y$ represent adoption-dependent burdens that remain when information quality is held fixed, such as additional maneuver effort, negotiation overhead, or onboard energy expenditure. In contrast, $q(r_L x + r_S y)$ represents the deterioration of the shared information state as adaptive-policy demand increases; that deterioration affects relative payoffs through $c_i q$. Thus, one pathway acts on utility at a given quality, while the other changes quality over time.

This separation must also be enforced in calibration. For example, if congestion-induced message delay is represented through the quality index, its resulting utility loss must not also be included in d_i. The direct-cost coefficient should be fitted from the residual dependence on focal adoption after controlling or experimentally holding q fixed, while r_i should be fitted to independently measured quality dynamics. A measured burden can enter both pathways only when separately identified consequences are modeled, such as energy expenditure at fixed quality and a separately measured loss of information freshness. Without such separation, the two effects may be confounded and should not be interpreted as independently calibrated mechanisms. Table 3 lists candidate measurements; it does not assign empirical values to the coefficients. Time and quality are dimensionless after normalization, and q is not identified directly with any single network or perception metric.

Unless otherwise stated, the linear-quality model satisfies

$$\kappa_i, c_i, d_i, r_i, \gamma > 0, \qquad a_i, b_i \in \mathbb{R}, \qquad 0 < q_{\max} \leq 1, \qquad i \in \{L, S\}. \tag{3.5}$$

Positive b_i describes cross-role complementarity, whereas negative b_i describes a reduction in the adaptive policy's relative payoff when its counterpart becomes more prevalent. Neither sign identifies the policy as aggressive or yielding. Improved quality is assumed to favor adaptation through $c_i > 0$. The equilibrium theorem below explicitly imposes its additional determinant and monotonicity conditions; signed interactions alone imply neither instability nor oscillation.

Table 3. Candidate observables for interpreting and calibrating the aggregate model.

Quantity	Model meaning	Possible observation or calibration contrast
$q, q_{\max}$	Effective information quality and unloaded upper bound	A prespecified index of task success, or a validated combination of message freshness and perception accuracy, normalized to [0,1]; low-load tests constrain its upper bound.
a_i, b_i	Baseline policy advantage and cross-role interaction	Adaptive-minus-fixed utility contrasts using fixed travel-time, comfort, risk, and social weights; vary the opponent share at matched quality and focal adoption.
c_i	Utility sensitivity to information quality	Utility changes when quality is varied while the two policy shares are controlled.
d_i	Direct adoption-dependent utility burden	Energy expenditure or maneuver effort converted to the common utility scale, conditional on focal adoption at fixed quality; exclude quality-mediated losses.
r_i	Adoption-induced quality depletion	Quality-rate changes under independently varied role-specific load at matched quality and recovery conditions.
γ	Quality recovery rate	Recovery trajectory toward $q_{\max}$ after adaptive-policy load is removed or controlled.
κ_i	Policy adjustment rate	Policy-share response to measured payoff differences after fixing the utility scale and time unit.

Note: Candidate measurements require an explicit quality-index definition and utility normalization; no mapping to a specific packet, latency, or perception measure is assumed. The multiplier ω scales both recovery and depletion and is a scenario response-speed factor, not an independently identifiable rate when γ and r_i are free.

3.3. Well-posedness and the interior equilibrium

3.3.1. Existence of solutions and positive invariance

Theorem 3.1. (Existence, uniqueness, and positive invariance) Under Eq (3.5), Eq (3.4) has a unique global solution for every initial state $(x_0, y_0, q_0) \in \Omega = [0,1]^2 \times [0, q_{\max}]$, and Ω is positively invariant. If $0 < x_0, y_0 < 1$, then $0 < x(t), y(t) < 1$ at every finite time. Moreover, information quality satisfies

$$q(t) \geq q_{\min} + (q_0 - q_{\min})e^{-(\gamma + r_L + r_S)t}, \quad q_{\min} = \frac{\gamma q_{\max}}{\gamma + r_L + r_S} > 0. \tag{3.6}$$

Proof. The right-hand side is polynomial and hence locally Lipschitz, so a unique local solution exists. On $x = 0,1$ or $y = 0,1$, the corresponding policy derivative vanishes. At the quality boundaries,

$\dot{q}|_{q=0} = \gamma q_{\max} > 0$ and $\dot{q}|_{q=q_{\max}} = -q_{\max}(r_L x + r_S y) \leq 0$. The vector field, therefore, does not point outside Ω. The solution remains in this compact set and extends to all $t \geq 0$.

For interior policy states, $\frac{d}{dt}\log\frac{x}{1-x} = \kappa_L g_L$, whose right-hand side is bounded on Ω. The log-odds cannot diverge in finite time, ruling out finite-time arrival of x at 0 or 1; the same argument applies to y. Finally, $\dot{q} \geq \gamma q_{\max} - (\gamma + r_L + r_S)q$. Comparison with the linear equation having the same initial value yields Eq (3.6). This bound also ensures that, even when $q_0 = 0$, $q(t) > 0$ for all $t > 0$. QED.

3.3.2. Existence and uniqueness of the interior equilibrium

Define

$$\begin{aligned} D &= d_L d_S - b_L b_S, \quad u_L = \frac{d_S a_L + b_L a_S}{D}, \quad v_L = \frac{d_S c_L + b_L c_S}{D}, \\ u_S &= \frac{b_S a_L + d_L a_S}{D}, \quad v_S = \frac{b_S c_L + d_L c_S}{D}, \\ A &= r_L v_L + r_S v_S, \qquad B = \gamma + r_L u_L + r_S u_S. \end{aligned} \tag{3.7}$$

Theorem 3.2. (Existence and uniqueness of the interior equilibrium) Under Eq (3.5), suppose that $D > 0$ and $A > 0$ in Eq (3.7). The only possible interior equilibrium of Eq (3.4) is

$$\begin{aligned} q^* &= \frac{-B + \sqrt{B^2 + 4A\gamma q_{\max}}}{2A}, \\ x^* = u_L + v_L q^*, &\qquad y^* = u_S + v_S q^*. \end{aligned} \tag{3.8}$$

This equilibrium exists if and only if

$$0 < x^* < 1, \qquad 0 < y^* < 1; \qquad \text{in this case, } q_{\min} < q^* < q_{\max}. \tag{3.9}$$

Proof. An interior policy equilibrium requires $d_L x - b_L y = a_L + c_L q$ and $-b_S x + d_S y = a_S + c_S q$. Since $D > 0$, the policy equations have the unique solution $x = u_L + v_L q$ and $y = u_S + v_S q$. Substitution into the quality equation gives $Aq^2 + Bq - \gamma q_{\max} = 0$. Under the explicit assumption $A > 0$, its roots have opposite signs, so Eq (3.8) is its unique positive root. For nonnegative cross-role coefficients, $D > 0$ already implies $v_L, v_S > 0$ and hence $A > 0$; signed coefficients require the additional check.

A positive quality root alone does not ensure feasible policy shares; Eq (3.9) must still be checked. If both shares lie strictly between 0 and 1, then $q^* = \gamma q_{\max}/(\gamma + r_L x^* + r_S y^*)$ directly gives $q_{\min} < q^* < q_{\max}$. Conversely, every interior equilibrium must satisfy the unique solution above and the policy feasibility conditions. The result follows. QED.

3.4. Global and local stability

3.4.1. A sufficient condition for global convergence

Let

$$w_L = \frac{r_L}{c_L}, \qquad w_S = \frac{r_S}{c_S},$$
$$\Delta = 4w_L w_S d_L d_S - (w_L b_L + w_S b_S)^2 > 0. \tag{3.10}$$

Theorem 3.3. (A sufficient condition for global stability) Under Eq (3.5), if a feasible interior equilibrium exists and Eq (3.10) holds, every solution with $(x_0, y_0, q_0) \in (0,1)^2 \times (0, q_{\max}]$ converges to it. The equilibrium is globally asymptotically stable relative to this domain, including signed cross-role coefficients.

Proof. Define the Bernoulli relative entropy and the Lyapunov function

$$\begin{aligned} \mathcal{B}(z^*\|z) &= z^* \log\frac{z^*}{z} + (1-z^*)\log\frac{1-z^*}{1-z}, \\ V(x,y,q) &= \frac{w_L}{\kappa_L}\mathcal{B}(x^*\|x) + \frac{w_S}{\kappa_S}\mathcal{B}(y^*\|y) \\ &\quad + q - q^* - q^* \log\frac{q}{q^*}. \end{aligned} \tag{3.11}$$

Each term in Eq (3.11) is nonnegative, and their sum vanishes only at the equilibrium. Write $X = x - x^*$, $Y = y - y^*$, and $Z = q - q^*$. Since $\partial_z \mathcal{B}(z^*\|z) = (z - z^*)/[z(1-z)]$, the derivative of the left-turning policy term is $w_L X(-d_L X + b_L Y + c_L Z)$, and that of the straight-going policy term is $w_S Y(-d_S Y + b_S X + c_S Z)$.

Let $G^* = \gamma + r_L x^* + r_S y^* = \gamma q_{\max}/q^*$. The quality equation can be rewritten as $\dot{q} = -G^* Z - q(r_L X + r_S Y)$. Hence, the derivative of the quality term is $(Z/q)\dot{q} = -(G^*/q)Z^2 - r_L XZ - r_S YZ$. Using $w_i c_i = r_i$, the cross terms between policies and quality cancel exactly, giving

$$\dot{V} = -w_L d_L X^2 - w_S d_S Y^2 + (w_L b_L + w_S b_S)XY - \frac{\gamma q_{\max}}{q q^*} Z^2. \tag{3.12}$$

The matrix of the policy quadratic form in Eq (3.12) is

$$S = \begin{pmatrix} w_L d_L & -(w_L b_L + w_S b_S)/2 \\ -(w_L b_L + w_S b_S)/2 & w_S d_S \end{pmatrix}. \tag{3.13}$$

The diagonal entries in Eq (3.13) are positive and its determinant is $\Delta/4 > 0$, so S is positive definite. Thus, $\dot{V} < 0$ at every nonequilibrium state. As x or y approaches 0 or 1, or as $q \downarrow 0$, $V \to \infty$. Together with $q \le q_{\max}$, this implies that the sublevel set of V containing the trajectory from any admissible initial state is compact, positively invariant, and bounded away from those boundaries. By LaSalle's invariance principle, the trajectory converges to the unique equilibrium at which $\dot{V} = 0$. The positive definiteness of V also establishes Lyapunov stability. QED.

The signed-interaction extension preserves this proof. Indeed, $(w_L b_L + w_S b_S)^2 - 4w_L w_S b_L b_S = (w_L b_L - w_S b_S)^2 \ge 0$, so $\Delta > 0$ implies $D > 0$. To see that it also implies $A > 0$, define $M = \begin{pmatrix} d_L & -b_L \\ -b_S & d_S \end{pmatrix}$, $W_d = \mathrm{diag}(w_L, w_S)$, $c = (c_L, c_S)^{\mathsf{T}}$, and $v = M^{-1}c$. The matrix in Eq (3.13) is $S = (W_d M + M^{\mathsf{T}} W_d)/2$. Since $W_d c = (r_L, r_S)^{\mathsf{T}}$, $A = (Mv)^{\mathsf{T}} W_d v = v^{\mathsf{T}} S v > 0$. Thus, the global condition supplies both assumptions of Theorem 3.2, but feasibility in Eq (3.9) must still be checked. Failure of $\Delta > 0$ leaves stability undecided.

The global result is restricted to states in which both policies are initially present. The invariant boundaries $x = 0,1$ or $y = 0,1$ are excluded. Moreover, asymptotic stability of the dynamical system

does not directly imply an evolutionarily stable strategy, optimal social welfare, or a traffic safety guarantee.

3.4.2. Local stability beyond the sufficient condition

Proposition 3.1. (Interior linearization) At any feasible interior equilibrium, set $\sigma_L = \kappa_L x^*(1-x^*)$, $\sigma_S = \kappa_S y^*(1-y^*)$, and $G^* = \gamma + r_L x^* + r_S y^*$. The Jacobian and its characteristic polynomial are

$$\begin{aligned} J^* &= \begin{pmatrix} -\sigma_L d_L & \sigma_L b_L & \sigma_L c_L \\ \sigma_S b_S & -\sigma_S d_S & \sigma_S c_S \\ -q^* r_L & -q^* r_S & -G^* \end{pmatrix}, \\ det(\lambda I - J^*) &= \lambda^3 + H_1\lambda^2 + H_2\lambda + H_3. \end{aligned} \tag{3.14}$$

$$\begin{aligned} H_1 &= \sigma_L d_L + \sigma_S d_S + G^*, \\ H_2 &= \sigma_L\sigma_S D + G^*(\sigma_L d_L + \sigma_S d_S) + q^*(\sigma_L c_L r_L + \sigma_S c_S r_S), \\ H_3 &= \sigma_L\sigma_S G^* D + \sigma_L\sigma_S q^* r_L(d_S c_L + b_L c_S) + \sigma_L\sigma_S q^* r_S(b_S c_L + d_L c_S). \end{aligned} \tag{3.15}$$

Direct differentiation gives Eq (3.14); the coefficients in Eq (3.15) follow from the trace, principal minors, and determinant. The Routh–Hurwitz criterion states that all eigenvalues have negative real parts if and only if

$$H_1 > 0, \qquad H_2 > 0, \qquad H_3 > 0, \qquad H_1 H_2 > H_3. \tag{3.16}$$

Hence, Eq (3.16) is an exact test of whether the linearization is Hurwitz and guarantees local exponential stability. A positive-real-part eigenvalue implies instability. A zero eigenvalue or a purely imaginary pair requires higher-order analysis. In particular, $H_1 H_2 = H_3$ with $H_1, H_2, H_3 > 0$ yields the factorization $(\lambda + H_1)(\lambda^2 + H_2)$ but establishes only a candidate Hopf point: Transversality and nonlinear nondegeneracy must also be verified. These tests remain applicable when $\Delta \le 0$ and, for a known feasible equilibrium, do not require $D > 0$ or $A > 0$.

At a boundary equilibrium, the general Jacobian has diagonal policy entries $\kappa_L[(1-2x)g_L - d_L x(1-x)]$ and $\kappa_S[(1-2y)g_S - d_S y(1-y)]$; its off-diagonal entries follow directly from Eq (3.4). These entries, rather than the interior simplification, are used for the edge equilibria in Section 4.6. At a policy corner $(\xi, \eta) \in \{0,1\}^2$, $q_{\xi\eta} = \gamma q_{\max}/(\gamma + r_L\xi + r_S\eta)$ and the eigenvalues reduce to $\kappa_L(1-2\xi)g_L$, $\kappa_S(1-2\eta)g_S$, and $-(\gamma + r_L\xi + r_S\eta)$, evaluated at that corner. Zero eigenvalues require nonlinear analysis.

3.4.3. Common nonlinear information benefits

Proposition 3.2. (Monotone nonlinear quality benefits) Replace $c_L q, c_S q$ in Eq (3.3) by $c_L h(q), c_S h(q)$, where the same continuously differentiable, strictly increasing function h is used for both roles. Normalize $h(0) = 0$ and $h(q_{\max}) = q_{\max}$. If a feasible interior equilibrium exists and $\Delta > 0$, it is globally asymptotically stable relative to the interior policy domain. Its quality solves

$$q^*[B + Ah(q^*)] = \gamma q_{\max},\ x^* = u_L + v_L h(q^*), y^* = u_S + v_S h(q^*). \tag{3.17}$$

There is at most one feasible solution because $\Delta > 0$ implies $D > 0, A > 0$, while $\gamma q_{\max}/q$ is strictly decreasing and $B + Ah(q)$ is strictly increasing for $q > 0$. Feasibility still requires $0 < x^*, y^* < 1$. For the stability proof, replace the quality term of Eq (3.11) by

$$V_h = \frac{w_L}{\kappa_L}\mathcal{B}(x^*\|x) + \frac{w_S}{\kappa_S}\mathcal{B}(y^*\|y) + \int_{q^*}^{q} \frac{h(s)-h(q^*)}{s}\, ds. \tag{3.18}$$

Strict monotonicity makes the integral positive away from q^* and divergent as $q \downarrow 0$. With $X = x - x^*$ and $Y = y - y^*$, differentiation cancels the policy–quality cross terms exactly and gives

$$\dot{V}_h = -(X \quad Y)S\begin{pmatrix} X \\ Y \end{pmatrix} - \frac{G^*}{q}(q - q^*)[h(q) - h(q^*)]. \tag{3.19}$$

Both terms are strictly dissipative away from equilibrium under $\Delta > 0$. The compact-sublevel and invariance argument from Theorem 3.3 proves convergence and stability. This extension covers smooth saturation and smooth threshold-like benefits, for example,

$$\begin{aligned} h_{\text{sat}}(q) &= \frac{q(K+q_{\max})}{K+q}, \\ h_{\text{thr}}(q) &= q_{\max}\frac{q^n}{K^n+q^n}\frac{K^n+q_{\max}^n}{q_{\max}^n},\ K > 0,\ n \geq 2. \end{aligned} \tag{3.20}$$

The normalization fixes the benefit at maximum quality, permitting comparisons that isolate response shape. A steep smooth threshold represents a rapid deterioration of relative benefit below an information-quality range, without prescribing a universal safety cutoff. Role-specific response functions and discontinuous hard thresholds are outside this proof because the cancellation or smoothness argument no longer follows. Numerical illustrations below distinguish this theorem from behavior outside its sufficient region.

3.5. Boundary invasion and policy invasion thresholds

At the all-fixed-policy equilibrium, define the role-specific invasion payoff indicators by

$$E_{00} = (0,0,q_{\max}), \qquad I_L = a_L + c_L q_{\max}, \qquad I_S = a_S + c_S q_{\max}. \tag{3.21}$$

Theorem 3.4. (Invasion conditions at the all-fixed-policy equilibrium) For Eq (3.21), if $I_L < 0$ and $I_S < 0$, then E_{00} is locally asymptotically stable. If at least one $I_i > 0$, then E_{00} is unstable, and a rare adaptive policy in that role has a positive linear invasion growth rate. If both indicators are nonpositive and at least one is zero, linearization alone cannot determine nonlinear stability.

Proof. At E_{00}, the Jacobian and its eigenvalues are

$$\begin{aligned} J(E_{00}) &= \begin{pmatrix} \kappa_L I_L & 0 & 0 \\ 0 & \kappa_S I_S & 0 \\ -q_{\max} r_L & -q_{\max} r_S & -\gamma \end{pmatrix}, \\ \lambda_1 = \kappa_L I_L, &\quad \lambda_2 = \kappa_S I_S, \quad \lambda_3 = -\gamma. \end{aligned} \tag{3.22}$$

The matrix in Eq (3.22) is lower triangular, so its eigenvalues are its diagonal entries. If all eigenvalues are negative, the equilibrium is locally asymptotically stable. If any invasion eigenvalue is positive, perturbations grow along the corresponding feasible policy direction and the equilibrium

is unstable. If there is a zero eigenvalue but no positive eigenvalue, further nonlinear analysis is required. QED.

For a specified dependence on external penetration p, the general indicator is $I_i(p) = a_i(p) + c_i(p)q_{\max}(p)$, and the corresponding linear invasion growth rate is $\lambda_i(p) = \kappa_i I_i(p)$. A zero of I_i is a candidate policy invasion threshold; its location, feasibility in $0 < p < 1$, and signs on both sides must be checked. Multiple roots and roots without a sign change are possible. The affine case is the special form

$$I_i(p) = \alpha_i p + \beta_i, \qquad p_{c,i} = -\frac{\beta_i}{\alpha_i}, \qquad \alpha_i \neq 0. \tag{3.23}$$

Equation (3.23) identifies a threshold in the mixed-traffic interval only when $0 < p_{c,i} < 1$. For $\alpha_i > 0$, invasion is possible above this value, whereas $\alpha_i < 0$ reverses the direction. Both negative indicators guarantee local exponential stability of the all-fixed-policy equilibrium; a positive indicator implies instability, whereas zero indicators require nonlinear analysis. A positive indicator describes the growth of a rare policy near this boundary and does not by itself prove interior coexistence from arbitrary initial states. These thresholds concern within-CAV policy adoption; they do not measure traffic-wave stability or safe deployment penetration.

3.6. Corollary on the information response timescale

To distinguish quality recovery capacity from information response speed, keep the policy equations unchanged and multiply the quality equation by a positive factor ω:

$$\dot{q} = \omega[\gamma(q_{\max} - q) - q(r_L x + r_S y)], \qquad \omega > 0. \tag{3.24}$$

Corollary 3.1. (Positive scaling of the quality equation) For each fixed $\omega > 0$, Eq (3.24) leaves all equilibria unchanged and preserves the global conclusion of Theorem 3.3.

Proof. The zeros of the quality equation are unchanged by multiplication by a positive constant. Keeping $w_i = r_i/c_i$ defined using the original coefficients, use

$$V_\omega(x, y, q) = \frac{w_L}{\kappa_L}\mathcal{B}(x^*\|x) + \frac{w_S}{\kappa_S}\mathcal{B}(y^*\|y) + \frac{1}{\omega}\left(q - q^* - q^*\log\frac{q}{q^*}\right). \tag{3.25}$$

The factor $1/\omega$ in Eq (3.25) cancels ω when the quality term is differentiated, leaving Eq (3.12) unchanged. The positive-definiteness and compact-sublevel arguments therefore apply for each fixed positive ω. QED. This corollary concerns the sufficient global stability region; outside that region, the local Jacobian must be checked with its third row multiplied by ω. Neither the corollary nor equilibrium preservation implies a monotonic change in convergence time. The adjustment rates κ_i likewise cancel from the Lyapunov criterion but remain in the Jacobian and transient dynamics.

The results concern the deterministic, delay-free, mean-field equations under the stated hypotheses. Section 4 tests their numerical consequences, the conservatism of the global sufficient condition, and the dependence of transient behavior on parameters. Empirical calibration and structural extensions are discussed in Section 5.

4. Numerical experiments and model verification

All experiments use synthetic, dimensionless parameters. The baseline calculations verify analytical equilibria and the Lyapunov identity; additional experiments examine parameter dependence, nonlinear information benefits, and dynamics outside the global sufficient region. The results concern policy adoption within each CAV role, with market penetration prescribed externally.

4.1. Experimental settings and numerical methods

Table 4 gives the baseline and asymmetric-role coefficients. The original linear-benefit experiments use Eq (3.4); the common-benefit extension uses Proposition 3.2. The original output intervals are 0.05 for general experiments and 0.5 for penetration scenarios. Additional signed-interaction, mechanism, and critical experiments are specified in Sections 4.6–4.9. All times are dimensionless evolutionary times.

Table 4. Parameters for the baseline and asymmetric-role experiments.

Parameter	Interpretation	Baseline	Asymmetric roles
a_L, a_S	Baseline payoff differences	(0.05, 0.05)	(0.08, 0.02)
b_L, b_S	Cross-role payoff coefficients	(0.20, 0.20)	(0.30, 0.10)
c_L, c_S	Information-benefit coefficients	(0.50, 0.50)	(0.60, 0.40)
d_L, d_S	Policy-adoption cost coefficients	(2.00, 2.00)	(2.00, 1.80)
r_L, r_S	Quality-depletion coefficients	(0.50, 0.50)	(0.70, 0.30)
κ_L, κ_S	Policy-update rates	(1.00, 1.00)	(1.00, 0.70)
γ	Information-recovery rate	1.00	1.00
$q_{\max}$	Quality level without additional load	1.00	1.00

The primary solver is the eighth-order DOP853 method, with relative and absolute tolerances of 10^{-10} and 10^{-12}, respectively, and a maximum internal step size of 0.5. The 75 baseline initial conditions are additionally checked using RK45, with relative and absolute tolerances of 10^{-12} and 10^{-14} and the same maximum step size of 0.5. These tolerances control local error estimates and do not provide an upper bound on the true error of an entire trajectory. The output sampling interval is also distinct from the solver's internal step size. No state clipping or projection is applied during integration.

Let $z = (x, y, q)^{\mathsf{T}}$. The sampled error and sampled convergence time are defined relative to the analytical equilibrium z^* of each scenario as

$$\begin{aligned} e_k &= \| z(t_k) - z^* \|_\infty, \qquad \varepsilon = 10^{-6}, \\ T_\varepsilon &= \min\{t_k : e_j \le \varepsilon \text{ for all } j = k, \dots, N\}. \end{aligned} \tag{4.1}$$

Equation (4.1) applies only to the observation interval $0 \le t_k \le T$; an empty set is reported as "not reached". This measure identifies the first sample at which the error meets the tolerance and

remains within it at every subsequent observed sample. It is not the exact continuous-time instant of permanent entry into the tolerance region. Convergence for the fixed-quality two-state control is evaluated using (x, y) only; the other experiments use all three states.

The computational environment is Python 3.12.14 with NumPy 2.3.5 and SciPy 1.17.0. The trajectory suite contains 90 original experiments, 137 signed-interaction, mechanism, and critical experiments, and 192 joint-parameter experiments, giving 419 primary integrations. The 80 independent solver checks comprise 75 baseline RK45 checks, three critical checks, and two joint-parameter Radau checks. Analytical equilibrium and eigenvalue evaluations are counted separately. Repeated deterministic integrations are not used to construct stochastic sampling confidence intervals.

All figures are drawn from the original trajectories. Points with errors below 10^{-9} are omitted from error plots, and policy shares below 10^{-6} are omitted from the penetration-dependent policy-evolution plot. These omissions affect visualization only: values are not replaced by a horizontal floor, and the calculation in Eq (4.1) is unchanged.

4.2. Convergence from multiple initial conditions and agreement with the analytical equilibrium

The baseline analytical equilibrium is $z^* = (0.25,0.25,0.80)$. The choices $x_0, y_0 \in \{0.01,0.10,0.50,0.90,0.99\}$ and $q_0 \in \{0.10,0.50,1.00\}$ yield 75 initial conditions, each integrated to $T = 100$. Figure 1 shows the policy shares of both roles, information quality, and equilibrium error; the main transient is displayed over $0 \le t \le 60$. Table 5 summarizes the corresponding numerical checks.

Table 5. Numerical consistency checks for the baseline model.

Check	Result
Initial conditions reaching $\varepsilon = 10^{-6}$	75 / 75
Range of sampled convergence times	27.40–35.95
Maximum terminal equilibrium error	2.220×10^{-16}
Maximum sampled difference between DOP853 and RK45	1.138×10^{-9}
Maximum violation of the state bounds	0
Maximum discrepancy in the Lyapunov derivative identity	8.882×10^{-16}

For the baseline coefficients, $w_i = r_i/c_i = 1$ and $\Delta = 15.84 > 0$. The Lyapunov derivative is checked independently by taking the inner product of its gradient with the vector field. All 75 initial conditions reach the prescribed tolerance, consistently with the global stability result; Table 5 reports the associated numerical discrepancies.

The solver comparison in Table 5 concerns the 75 baseline cases. The interior convergence theorem applies to $0 < x_0, y_0 < 1$ and $q_0 > 0$; policy-share boundaries at 0 and 1 are invariant and are treated separately.

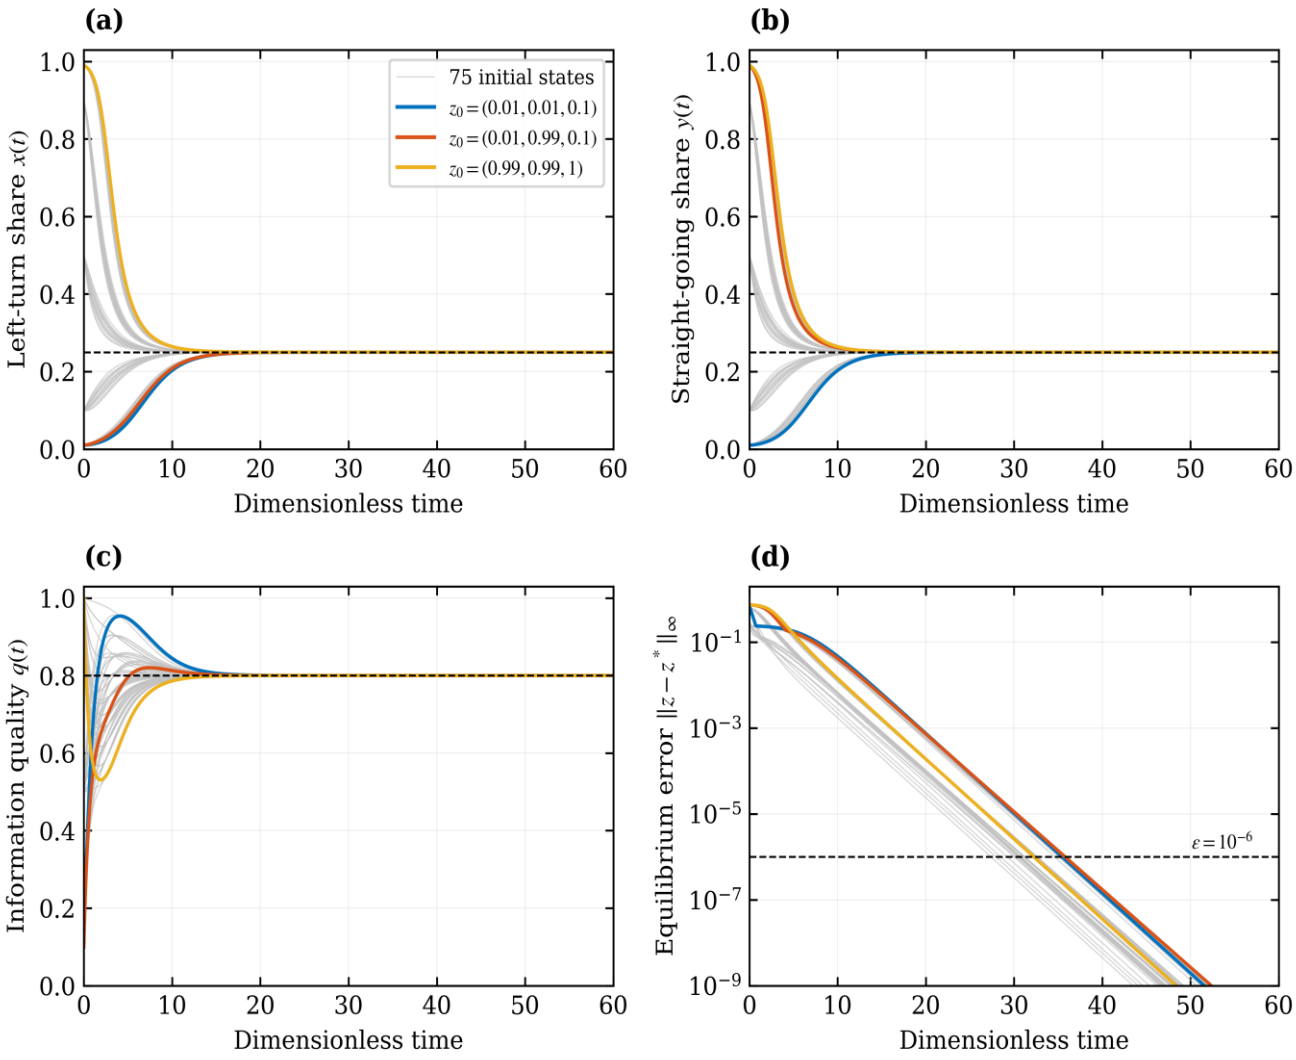


Figure 1. Baseline results for 75 initial conditions. (a) Left-turn policy share; (b) straight-going policy share; (c) effective information quality; (d) maximum-norm error relative to the analytical equilibrium. Gray curves represent all initial conditions, and the colored curves identify the same three representative initial conditions across panels. Horizontal dashed lines in (a–c) indicate the corresponding equilibrium values; the dashed line in (d) marks the 10^{-6} error tolerance. Integration continues to $T = 100$, whereas the displayed interval ends at 60. Points below 10^{-9} are omitted in (d).

4.3. Asymmetry between left-turn and straight-going roles

Using the asymmetric coefficients in Table 4, the initial state $(x_0, y_0, q_0) = (0.10, 0.90, 0.20)$ is integrated to $T = 150$. Figure 2(a) shows the two roles approaching different policy shares. Figure 2(b) projects the same three-dimensional trajectory onto the (x, y) plane. The equilibrium and convergence results are given in Table 6.

The left-turn share is larger in this specified asymmetric case. Because several coefficients change together, the example illustrates a combined role difference. Section 4.7 isolates payoff asymmetry

while holding the other coefficients fixed and then compares information mechanisms across three parameter backgrounds.

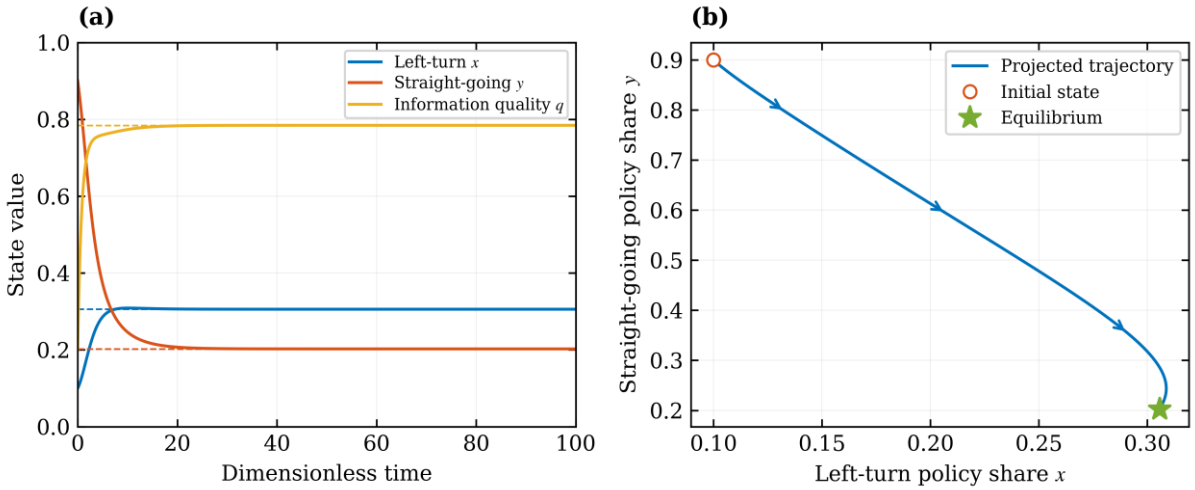


Figure 2. Asymmetric-role scenario. (a) Time histories of the left-turn share, straight-going share, and information quality, with dashed lines indicating their analytical equilibria. (b) Projection of the same three-dimensional trajectory onto the policy plane; the open circle marks the initial state, the star marks the equilibrium, and the arrows indicate increasing time. Panel (a) displays the interval up to 100, although integration continues to 150. Panel (b) is neither a complete phase-plane vector field nor a basin of attraction.

Table 6. Equilibrium and convergence results for asymmetric roles.

Quantity	Result
Equilibrium left-turn share x^*	0.305709
Equilibrium straight-going share y^*	0.202425
Equilibrium quality q^*	0.784484
Sampled convergence time T_ε	59.05
Maximum terminal equilibrium error	3.192×10^{-15}
Stability-condition margin Δ	12.419375

4.4. Information-recovery capacity, response speed, and endogenous feedback

4.4.1. Information-recovery capacity

Keeping the other baseline coefficients unchanged, we set $\gamma \in \{0.25,1,4\}$ and integrate each case from $(0.10,0.90,0.20)$ to $T = 150$. Figure 3 compares the policy and quality trajectories, and Table 7 reports the long-run equilibria and sampled convergence times.

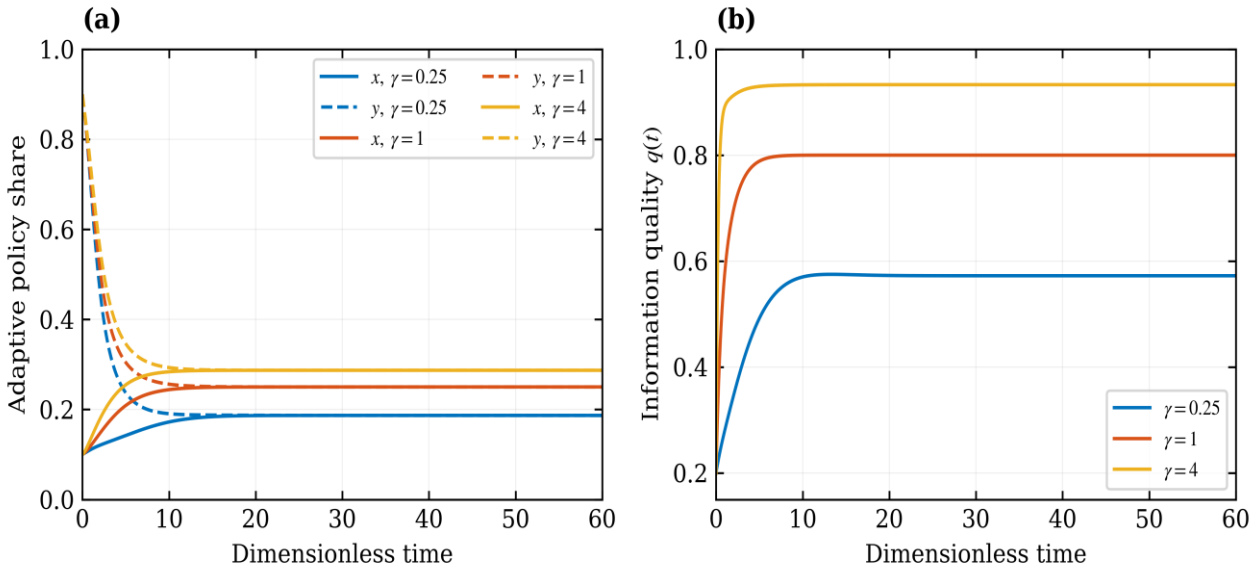


Figure 3. Results for different information-recovery rates. (a) Left-turn shares x are shown by solid curves and straight-going shares y by dashed curves; matching colors indicate the same γ. (b) Quality trajectories for the three recovery rates. All cases use the same initial state and are integrated to 150, with the displayed interval ending at 60. Changing γ alone changes both the equilibrium and the transient.

Table 7. Equilibrium and convergence results for different information-recovery rates.

γ	$x^* = y^*$	q^*	T_ε
0.25	0.186772	0.572381	37.50
1.00	0.250000	0.800000	31.15
4.00	0.286962	0.933062	30.20

Increasing γ from 0.25 to 4 raises the equilibrium quality from approximately 0.5724 to 0.9331 and the equilibrium policy share from approximately 0.1868 to 0.2870. Thus, recovery capacity changes the long-run balance between quality restoration and adoption-induced load. More generally, along a feasible equilibrium branch of Theorem 3.2, implicit differentiation gives $dq^*/d\gamma = (q_{\max} - q^*)/(G^* + Aq^*) > 0$. Under nonnegative b_i and $D > 0$, $v_i > 0$, so $dx^*/d\gamma = v_L dq^*/d\gamma > 0$ and likewise for y^*. These equilibrium derivatives do not determine convergence time.

4.4.2. Information-quality response timescale

To distinguish recovery capacity from response speed, we use the quality-timescale model introduced in Section 3.

Equation (3.24) scales the recovery and depletion rates together while preserving the equilibrium $(0.25,0.25,0.80)$. We set $\omega \in \{0.1,1,10\}$, use the common initial state $(0.10,0.90,0.80)$, and

integrate to $T = 200$. Figures 4(a) and 4(b) compare quality and full-state equilibrium error; the results are summarized in Table 8.

Table 8. Information-response timescales and the fixed-quality control.

Setting	Sampled minimum q	T_ε	Observation horizon T
Endogenous quality, $\omega = 0.1$	0.771231	68.45	200
Endogenous quality, $\omega = 1$	0.708383	31.20	200
Endogenous quality, $\omega = 10$	0.673614	31.65	200
Fixed quality, $q = 0.8$	0.800000	34.95	100

Note: Convergence in the fixed-quality case is evaluated using (x, y) only; the other cases use (x, y, q). Minimum quality is the minimum over observed samples; a continuous-time minimum is not computed.

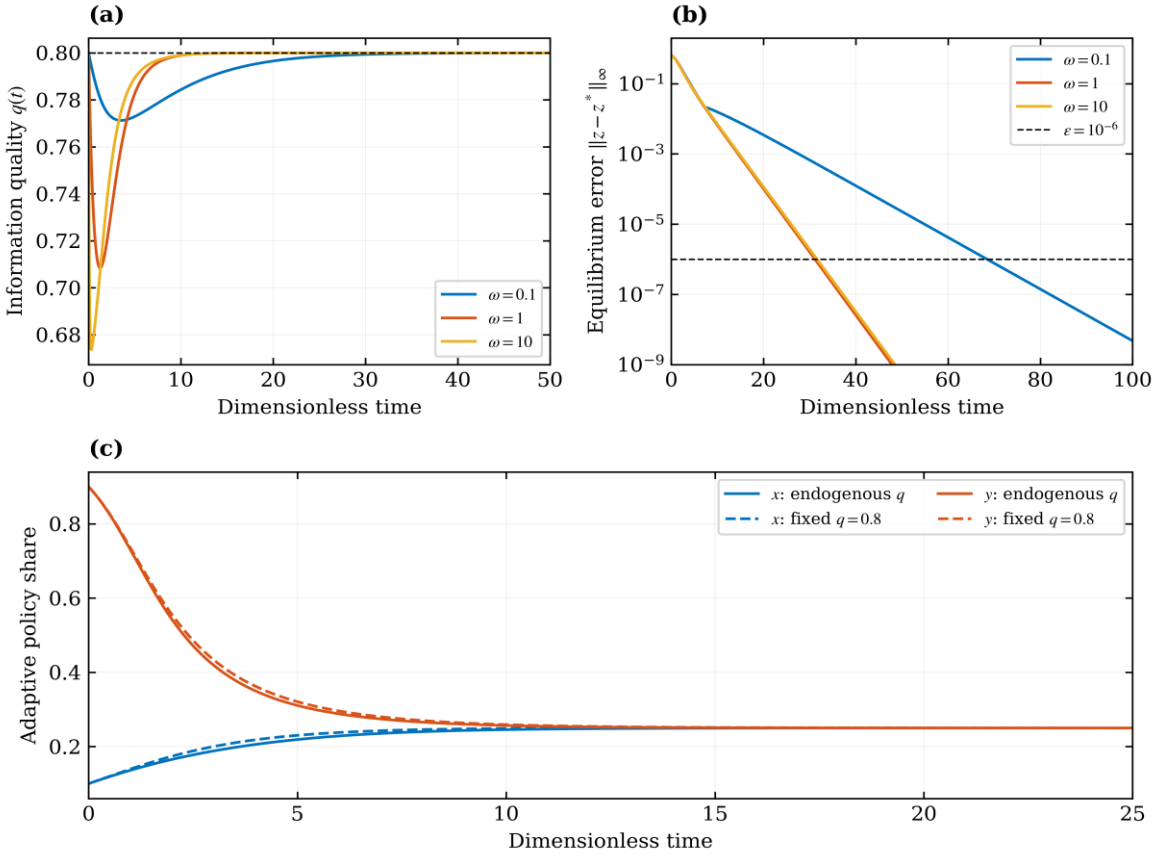


Figure 4. Information-response speeds and the fixed-quality control. (a) Quality transients for $\omega = 0.1,1,10$; the dashed line indicates their common equilibrium $q^* = 0.80$. (b) Corresponding full-state equilibrium errors; the dashed line marks the 10^{-6} tolerance, and plotting points below 10^{-9} are omitted. (c) Policy trajectories under fixed and endogenous quality; solid curves indicate endogenous quality, dashed curves indicate fixed $q = 0.80$, and the two roles are distinguished by color. Panels (a) and (b) display intervals ending at 50 and 100, respectively, with both integrated to 200. Panel (c) displays the interval ending at 25, with integration continuing to 100.

The three original values give a counterexample to a universally decreasing sampled convergence time: T_ε falls from 68.45 to 31.20 and then rises to 31.65. They do not describe a general response curve. At their common initial state, $\dot{q}(0) = -0.2\omega$, so greater response speed also accelerates the initial quality decline. Section 4.7 examines the time-scale dependence on a 25-point grid in three backgrounds.

4.4.3. Fixed versus endogenous quality

We also compare the two-state model with fixed $q \equiv 0.80$ against the baseline endogenous-quality model. Both use the initial policy shares $(0.10,0.90)$; the endogenous model additionally uses $q_0 = 0.80$. Both are integrated to $T = 100$. Figure 4(c) shows the resulting differences in policy transients.

The two quality settings share the policy equilibrium $(0.25,0.25)$, but the maximum sampled difference between their policy trajectories is 0.014646. Policy-only convergence times are 34.95 and 31.20 for fixed and endogenous quality, respectively. In this endogenous-quality case, policy-only and full-state convergence times both equal 31.20, illustrating how quality feedback modifies policy adjustment even when its endpoint is preserved.

4.5. Policy invasion and near-critical convergence

4.5.1. A synthetic payoff-parameter family

This subsection uses a separate symmetric payoff-parameter family:

$$a(p) = 0.9p - 0.55, \qquad b(p) = 0.4p, \qquad c(p) = 0.2(1-p). \tag{4.2}$$

Both roles use $d = 2$ and $r = 0.5$, with $\kappa_L = \kappa_S = \gamma = q_{\max} = 1$. Equation (4.2) defines an illustrative parameter family, not a traffic calibration. In particular, $c(p)$ represents the marginal payoff sensitivity to quality, not the amount of information supplied. Its decrease with p is a scenario assumption. Near $(0,0,1)$, the linear invasion growth rate is

$$I(p) = a(p) + c(p) = 0.7p - 0.35, \qquad p_c = 0.50. \tag{4.3}$$

The symmetric interior equilibrium $x^* = y^* = s$ satisfies

$$\begin{gathered}(2 - 0.4p)s^2 + (2.55 - 1.3p)s - (0.7p - 0.35) = 0, \\ q^* = \frac{1}{1+s}.\end{gathered} \tag{4.4}$$

For $0.50 < p < 1$, Eq (4.4) has a feasible positive root. For $p \le 0.50$, the extinction equilibrium used for comparison is $(0,0,1)$. We choose $p \in \{0.30,0.45,0.49,0.51,0.55,0.70\}$ and integrate from the common initial state $(10^{-4}, 10^{-4}, 1)$ to $T = 2000$. Figure 5 presents the six numerical solutions together with the analytical curves obtained from Eqs (4.3) and (4.4). The 181 analytical parameter points are not additional ODE simulations.

For this particular parameter family, nonlinear extinction can also be established directly for $0 \le p \le 0.50$. Let $W = -\log(1-x) - \log(1-y)$. For interior policy states and $0 \le q \le 1$, the inequality $2xy \le x^2 + y^2$ gives

$$\begin{aligned}\dot{W} &= I(p)(x+y) + c(p)(q-1)(x+y) - 2(x^2+y^2) + 0.8p\,xy \\ &\le I(p)(x+y) - (2-0.4p)(x^2+y^2).\end{aligned} \tag{4.5}$$

For $p \le 0.50$, Eq (4.5) is strictly negative whenever the policy state is nonzero. Its sublevel sets keep x, y away from 1 and imply $\int_0^\infty (x^2+y^2)\,dt < \infty$. Bounded derivatives make x^2+y^2 uniformly continuous, so $x, y \to 0$. The stable quality equation then gives $q \to 1$. This includes the critical point, whose slow algebraic decay is analyzed below and tested numerically in Section 4.8.

4.5.2. Policy evolution and finite-time results

Table 9 reports the analytical equilibria, finite-time endpoints, and sampled convergence times for the six scenarios.

Figure 5(a–b) and Table 9 show that adaptive policies decay for $p < 0.50$, whereas rare adaptive policies grow and approach a positive coexistence equilibrium for $p > 0.50$. A positive invasion rate does not lead to complete replacement of the fixed policy: For example, the equilibrium share at $p = 0.70$ is approximately 0.078846. Policy extinction denotes the disappearance of that policy within the CAV population, not the withdrawal of CAVs from the market.

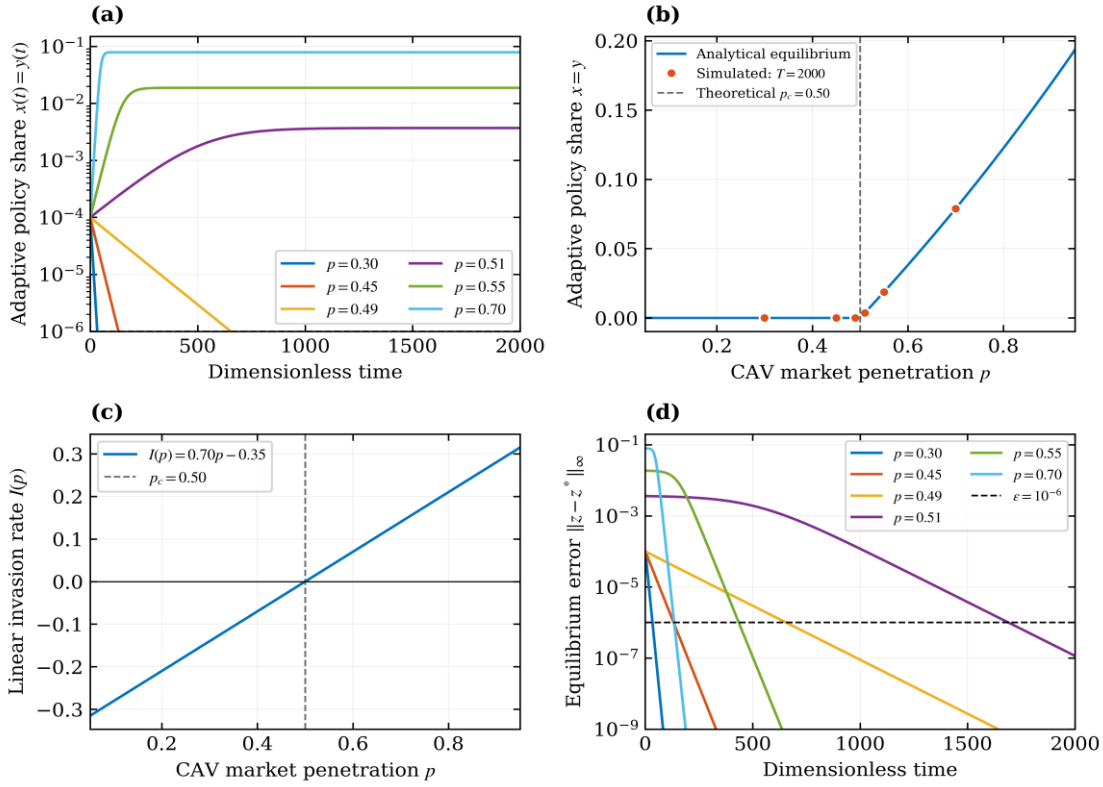


Figure 5. The original six off-critical scenarios for the synthetic penetration family. (a) Policy shares with values below 10^{-6} omitted from the logarithmic plot; (b) analytical equilibrium branch and six endpoints at $T = 2000$; (c) linear invasion rate; (d) full-state equilibrium errors, with values below 10^{-9} omitted. The vertical reference is this family's policy invasion threshold, $p_c = 0.50$.

Table 9. Numerical results for the six penetration scenarios.

p	$I(p)$	Analytical $x^* = y^*$	$x(2000) = y(2000)$	$q(2000)$	T_ε
0.30	−0.140	0	$< 10^{-6}$	1.000000000	34.0
0.45	−0.035	0	$< 10^{-6}$	1.000000000	132.5
0.49	−0.007	0	8.09567×10^{-11}	1.000000000	655.5
0.51	0.007	0.003696586	0.003696472	0.996317143	1689.5
0.55	0.035	0.018733157	0.018733157	0.981611321	436.0
0.70	0.140	0.078845922	0.078845922	0.926916421	135.0

Note: All six scenarios reach tolerance in Eq (4.1) within the observation interval. Values of $q(2000)$ are rounded to nine decimal places; the unrounded value for $p = 0.49$ is approximately 0.999999999918.

4.5.3. Near-critical convergence and interpretation of errors

The slow transients near the synthetic threshold can be related directly to the linearization. Let $\delta = p - 0.50$. At extinction, the policy eigenvalues are both 0.7δ, whereas the quality eigenvalue is -1. In the symmetric invariant subspace $x = y = s$, a local center-manifold reduction around $(s, q, \delta) = (0,1,0)$ yields

$$\begin{aligned} q &= 1 - s + O(s^2 + |\delta|s), \\ \dot{s} &= 0.7\delta s - 1.9s^2 + O(|\delta|s^2 + s^3), \\ s^* &= \frac{0.7}{1.9}\delta + O(\delta^2) \qquad (\delta > 0). \end{aligned} \tag{4.6}$$

For $\delta > 0$, the slow symmetric eigenvalue at coexistence is $-0.7\delta + O(\delta^2)$; the antisymmetric policy eigenvalue is $-(2.2)(0.7/1.9)\delta + O(\delta^2)$, and the remaining eigenvalue tends to -1. Consequently, the local exponential relaxation scale on either side, and the critical symmetric trajectory, satisfy

$$\tau \sim \frac{1}{0.7|p-0.50|} \quad (p \neq 0.50), \qquad s(t) \sim \frac{1}{1.9t}, \quad 1 - q(t) \sim s(t) \quad (p = 0.50). \tag{4.7}$$

To obtain Eq (4.6), write $q = 1 + \eta$. The symmetric quality equation is $\dot{\eta} = -\eta - s - \eta s$, so its center manifold has $\eta = -s + O(s^2 + |\delta|s)$. Substituting this relation into the policy equation gives the quadratic coefficient -1.9. At criticality, $d(1/s)/dt = 1.9 + o(1)$ along a positive symmetric trajectory approaching extinction, which establishes the algebraic decay in Eq (4.7). The full critical system has two zero eigenvalues; this symmetric reduction is not a classification of a generic one-dimensional bifurcation.

The scale τ is defined by the dominant local eigenvalue, whereas T_ε is a finite-time tolerance measure that also depends on the initial distance, nonlinear evolution, and sampling. Their asymptotic behavior need not coincide for fixed initial conditions and a fixed tolerance. The critical solution can therefore remain above the numerical tolerance over a long finite observation interval while still converging algebraically. The original off-critical cases at $p = 0.49$ and 0.51 have sampled convergence times of 655.5 and 1689.5, respectively; the unequal values are consistent with their different paths and distances to their own equilibria. Analytical equilibria and finite-time endpoints are reported separately.

4.6. Signed interactions and the limits of the sufficient condition

The additional trajectory suite consists of 44 signed-interaction runs, 75 response-speed runs, 7 common-benefit runs, 4 low-quality controls, and 7 extended near-critical runs. The first 130 use DOP853 with $\mathrm{rtol} = 10^{-10}$, $\mathrm{atol} = 10^{-12}$, maximum step 0.5, and output spacing 0.05. Signed cases use $T = 400$; nonlinear cases use $T = 200$. Response-speed horizons are $\max(200, 20/\omega)$ rounded to the nearest output interval. The critical solver and horizons are specified in Section 4.8. No clipping or projection is applied.

With all other baseline coefficients fixed, b_L, b_S are independently varied over 121 equally spaced values in $[-4,4]$. When $D \neq 0$, all real roots of the reduced quadratic are checked for feasible x, y, q; the single-positive-root formula is not assumed for $D < 0$. The 18 grid points with $D = 0$ are excluded from this elimination. Figure 6(a) classifies 5959 regular grid points as stable with $\Delta > 0$, 88 as locally stable despite $\Delta \leq 0$, 900 as having an unstable interior equilibrium, and 7676 as having no feasible interior equilibrium. No multiple regular feasible interiors occur on this grid. These are interior classifications, not classifications of all boundary attractors. For numerical classification, values satisfying $|\Delta| \leq 10^{-10}$ were treated as lying on the boundary $\Delta = 0$.

Four coefficient pairs are tested from 11 initial states: $(x_0, y_0) \in \{0.1, 0.5, 0.9\}^2$, plus $(0.500001, 0.499999)$ and its role reversal, with $q_0 = 0.5$. Table 10 gives the outcomes. The first three cases reach their interior equilibrium from all tested starts. In particular, $(b_L, b_S) = (4, 0.1)$ has $\Delta = -0.81$ but eigenvalues -1.403008 and $-0.423258 \pm 0.092234\,\mathrm{i}$; the stable trajectories in Figure 6(b) therefore lie outside the sufficient global certificate.

For $b_L = b_S = -3$, the interior point $(0.100840, 0.100840, 0.908397)$ has a positive eigenvalue 0.090671. The eight asymmetric starts split equally between $(0.247472, 0, 0.889889)$ and $(0, 0.247472, 0.889889)$. The full Jacobian at either edge has eigenvalues -0.432387, -1.063808, and -0.247472, confirming local stability. The three exactly symmetric starts remain in $x = y$ and approach the interior saddle. Figure 6(c–d) thus expose an instability that symmetric starts conceal. No sustained periodic motion is observed in these finite-window trajectories; the tests do not exclude periodic attractors elsewhere.

Table 10. Signed-interaction cases with the remaining baseline coefficients fixed.

Case / (b_L, b_S)	Interior (x^*, y^*, q^*)		Δ	max Reλ	Observed limits / 11 starts
B (0.2, 0.2)	(0.250000,	0.250000, 0.800000)	15.84	-0.412500	11 interior
S (−0.5, −0.5)	(0.188307,	0.188307, 0.841534)	15.00	-0.229271	11 interior
O (4, 0.1)	(0.659870,	0.230955, 0.691844)	-0.81	-0.423258	11 interior
C (−3, −3)	(0.100840,	0.100840, 0.908397)	-20.00	0.090671	3 symmetric to saddle; 4 to each edge

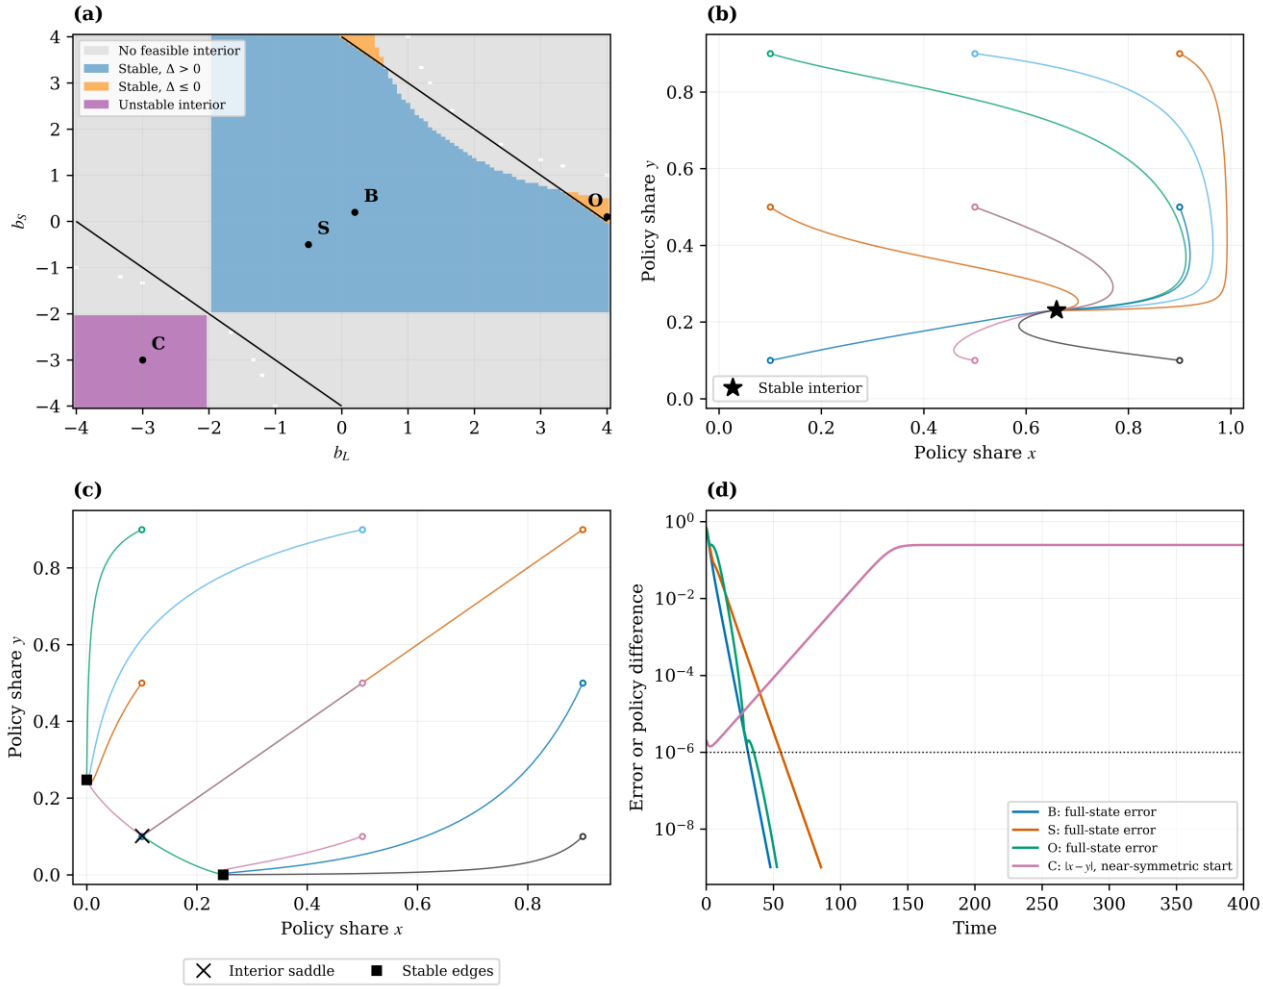


Figure 6. Signed interactions and stability. (a) Interior classification on the 121×121 grid; black lines mark $\Delta = 0$, and 18 points with $D = 0$ are excluded. B, S, O, and C refer to Table 10. (b,c) Policy-plane projections of 11 trajectories for O and C; the star is a stable interior equilibrium, and the cross is an interior saddle. (d) Full-state errors for B, S, and O from $(0.1,0.9,0.5)$, and $|x - y|$ from a nearly symmetric competitive start. The dotted line is 10^{-6}. Integrations continue to $T = 400$; the projections are not complete phase-plane vector fields.

4.7. Parameter dependence and nonlinear quality benefits

To isolate payoff asymmetry, set $a_L = 0.05 + \delta_a$ and $a_S = 0.05 - \delta_a$, retaining all other baseline coefficients. Across 41 equally spaced values $\delta_a \in [-0.04,0.04]$, $q^* = 0.8$ while $x^* - y^*$ varies from -0.036364 to 0.036364 (Figure 7(a)). This controlled contrast reverses the role ordering without simultaneously changing costs or resource loads.

The recovery and response-speed scans use three backgrounds: the baseline; the baseline with $b_L = b_S = -0.5$; and the asymmetric coefficients of Table 4. For 25 logarithmically spaced $\gamma \in [0.1,10]$ per background, equilibrium quality increases from 0.412573 to 0.971107, from 0.468115

to 0.978878, and from 0.393142 to 0.967946 (Figure 7(b)). These 75 equilibrium evaluations are algebraic calculations, not additional trajectories.

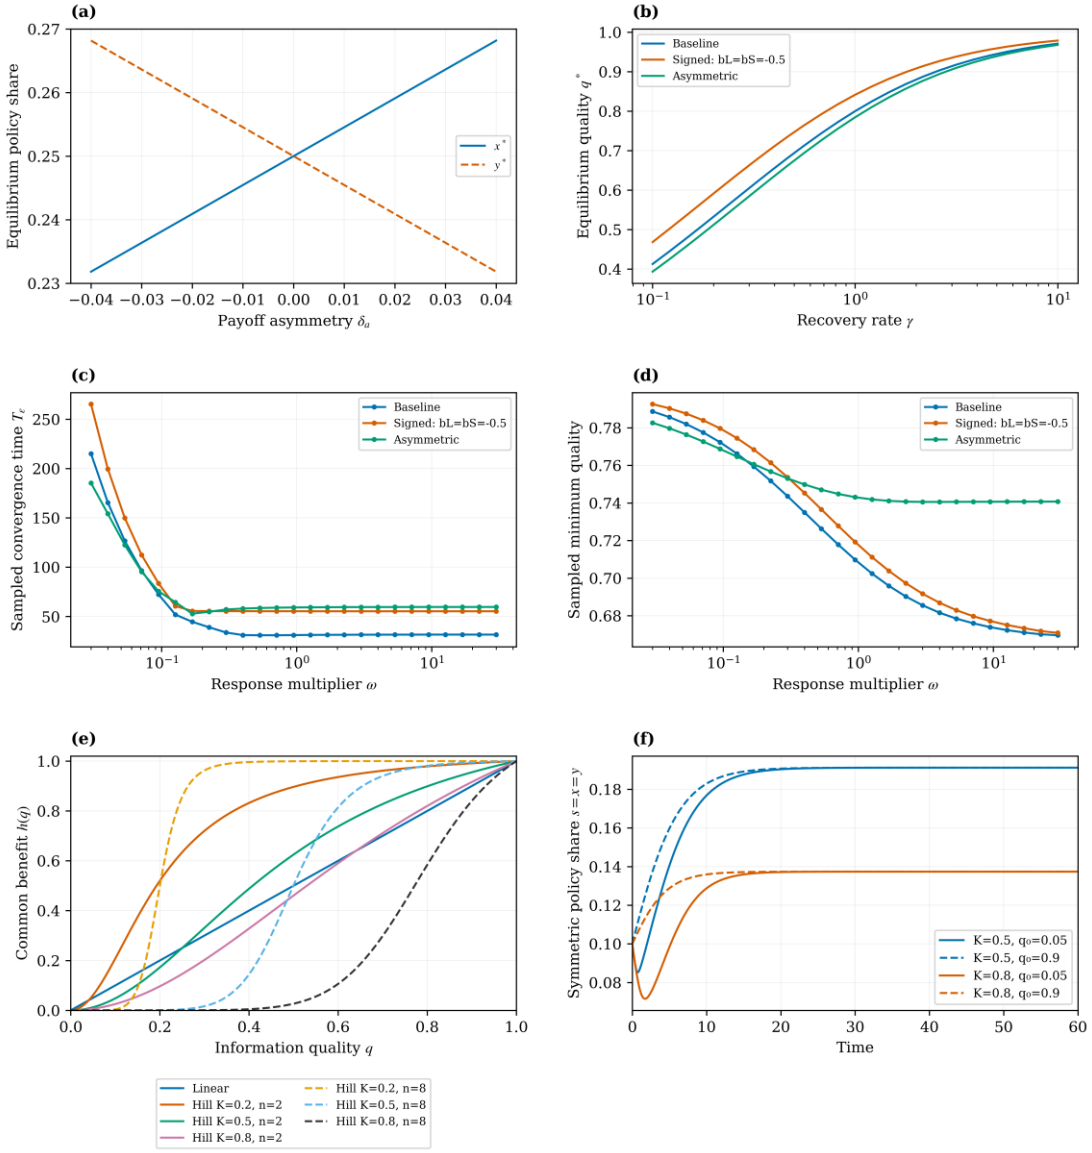


Figure 7. Controlled mechanism scans. (a) Isolated payoff asymmetry; (b) equilibrium quality for 25 recovery rates in each of three backgrounds; (c,d) sampled convergence time and minimum quality for 25 response multipliers per background; (e) linear and normalized Hill benefits; (f) low-quality controls with $a_L = a_S = -0.15$, $n = 8$, and $x_0 = y_0 = 0.1$. Solid and dashed curves in (f) use $q_0 = 0.05$ and 0.9. Remaining coefficients are baseline values. Recovery capacity and response speed are varied separately; all quantities are dimensionless.

Separately, 25 logarithmically spaced $\omega \in [0.03,30]$ are integrated per background from $(0.1,0.9,0.8)$. Figure 7(c–d) show that invariant equilibrium locations coexist with changing convergence times and minimum qualities. For the asymmetric background, the smallest sampled T_ε

is 52.85 at $\omega \approx 0.168702$, followed by longer times at larger ω. This establishes a nonmonotonic response in the specified family without identifying a universal optimal speed.

The common nonlinear benefit is the normalized Hill function in Eq (3.20), with $q_{\max} = 1$, $K \in \{0.2,0.5,0.8\}$, and $n \in \{2,8\}$, compared with $h(q) = q$. All responses share $h(0) = 0$ and $h(1) = 1$ (Figure 7(e)). With baseline coefficients, the seven trajectories from $(0.1,0.9,0.8)$ converge by $T = 200$. The equilibrium policy share ranges from 0.210622 to 0.305550, equilibrium quality from 0.765961 to 0.826022, and sampled convergence time from 29.75 to 34.55. Smooth nonlinear response shapes therefore change adoption levels while satisfying Proposition 3.2.

A further low-quality control sets $a_L = a_S = -0.15$, $n = 8$, and $K = 0.5$ or 0.8, with $x_0 = y_0 = 0.1$ and $q_0 = 0.05$ or 0.9. For low q_0, both initial payoff differences are approximately -0.33; shares first fall to 0.085264 or 0.071561, then recover toward 0.191183 or 0.137364 as quality improves. The corresponding high-quality initial gains are 0.167439 and 0.090139. Figure 7(f) shows this reversal in initial policy advantage. These smooth threshold-sensitive controls do not implement a discontinuous fail-safe switch or identify a calibrated communication threshold.

4.8. Policy invasion thresholds and critical decay

The policy invasion threshold is not fixed by the general model. Holding $b(p), c(p)$ as in Eq (4.2), replacing $a(p) = 0.9p - 0.55$ by $a(p) = 0.9p - 0.48$ gives $I(p) = 0.7(p - 0.4)$ and shifts the policy invasion threshold to 0.4 (Figure 8(a)). Other functional dependencies can remove a zero or produce multiple roots. This comparison treats penetration analysis as a template whose coefficients require calibration.

Seven full three-state trajectories use $p \in \{0.490,0.495,0.499,0.500,0.501,0.505,0.510\}$ and $(x_0, y_0, q_0) = (10^{-4}, 10^{-4}, 1)$. LSODA uses $\text{rtol} = 10^{-10}$, $\text{atol} = 10^{-13}$, and maximum step 100. Output spacing is 0.5 through $t = 2000$ and 100 thereafter, including the final endpoint. Noncritical horizons are $100\lceil 40/[100(0.7)|p - 0.5|]\rceil$; the critical horizon is 600,000. Table 11 distinguishes finite-horizon nonconvergence from convergence during the longer observation window.

Table 11. Near-critical convergence with $\varepsilon = 10^{-6}$.

p	$e(2000)$	T_ε by 2000	Extended T	T_ε extended
0.490	1.096×10^{-10}	655.5	5800	655.5
0.495	8.680×10^{-8}	1302.0	11,500	1302.0
0.499	2.049×10^{-5}	Not reached	57,200	6300.0
0.500	7.247×10^{-5}	Not reached	600,000	521,100.0
0.501	1.468×10^{-4}	Not reached	57,200	9900.0
0.505	2.911×10^{-5}	Not reached	11,500	3000.0
0.510	1.147×10^{-7}	1689.5	5800	1689.5

Note: Values after $t = 2000$ use a 100-unit output interval and are sampled convergence times, not exact continuous-time crossings. The independent LSODA endpoints can differ from the original DOP853 values in Table 9 at the solver-discrepancy scale.

At $p = 0.50$, $e(2000) = 7.247 \times 10^{-5}$, but the extended sampled convergence time is 521,100. At $t = 600{,}000$, the error is 8.696×10^{-7} and $ts(t) = 0.521741$, within 0.87% of the asymptotic value $1/1.9 = 0.526316$ (Figure 8(d)). A 200-point two-sided eigenvalue scan supports the spectral scaling in Eq (4.7). The tolerance crossing in Figure 8(c) remains a separate finite-observation quantity.

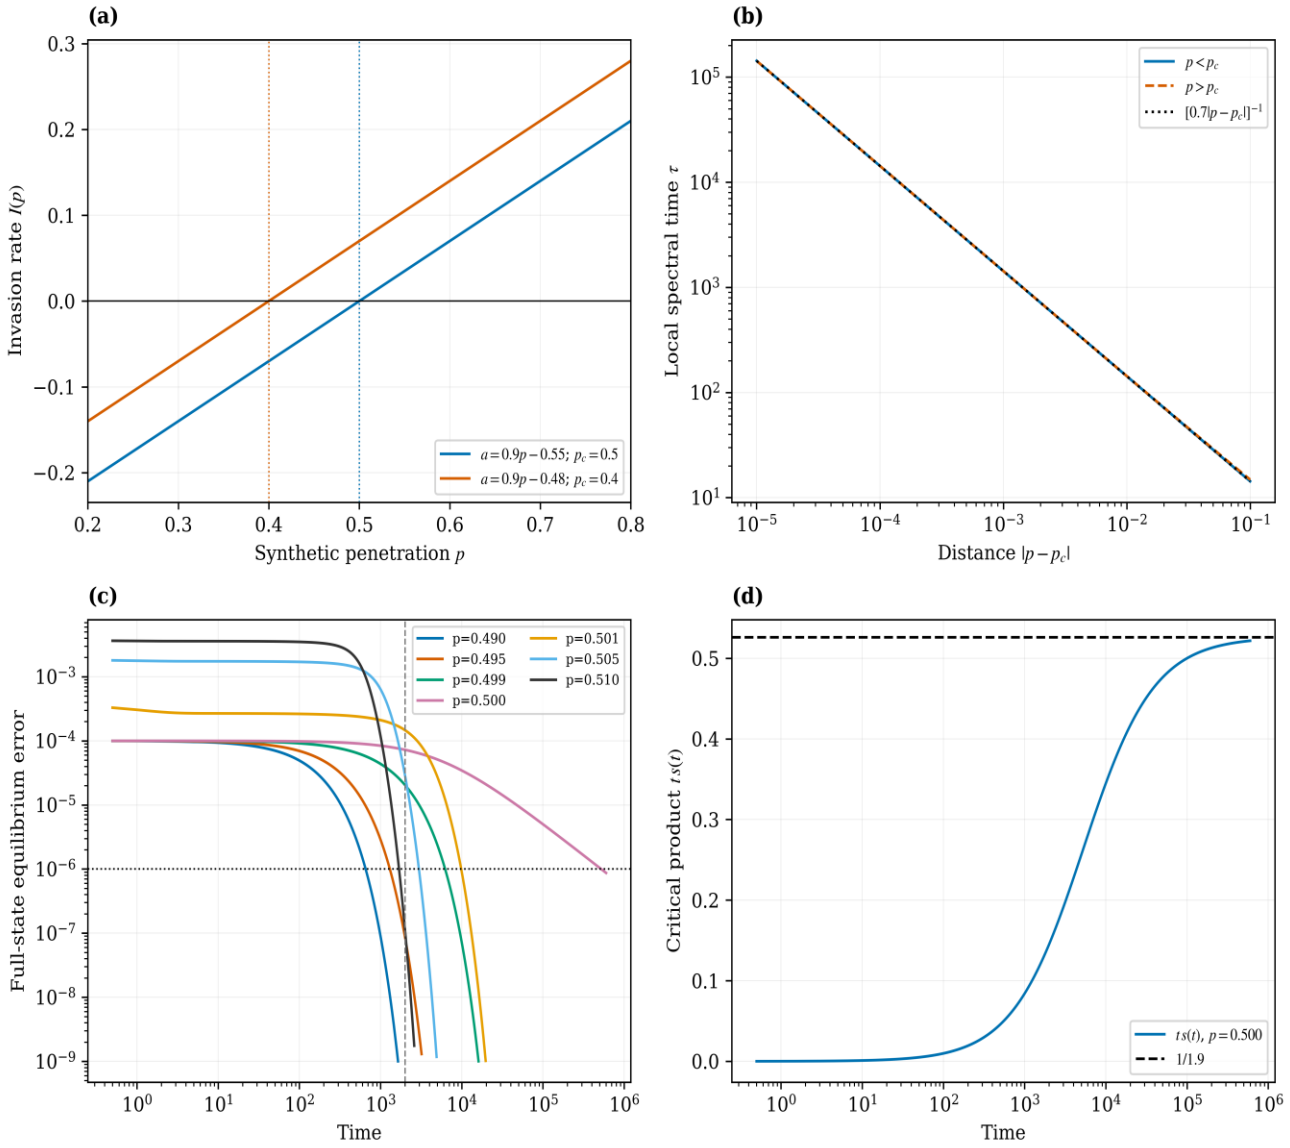


Figure 8. Policy invasion and critical convergence. (a) Two synthetic intercept choices give policy invasion thresholds of 0.5 and 0.4; (b) local spectral relaxation times on both sides compared with $[0.7|p - p_c|]^{-1}$; (c) full-state errors from symmetric initial states, with references at 10^{-6} and $T = 2000$; (d) the critical product $ts(t)$ approaches $1/1.9$. Trajectories have the parameter-dependent horizons and output intervals in the text. The critical asymptotics concern the invariant subspace $x = y$.

The complete critical trajectory was independently checked with Radau ($\text{rtol} = 10^{-11}$, $\text{atol} = 10^{-14}$, maximum step 1000); the largest sampled difference from LSODA was 8.375×10^{-11}, and the sampled convergence time was unchanged. DOP853 checks for $p = 0.499, 0.500$ through $T = 2000$,

using maximum step 0.5 and tolerances $10^{-10}, 10^{-13}$, differed by less than 7.30×10^{-11}. These are solver-comparison observations, not formal trajectory-error bounds.

4.9. Systematic joint-parameter sampling

To examine simultaneous changes in strategic coupling and information dynamics, we use all combinations of $b_L, b_S \in \{-3, -0.5, 0.2, 4\}$, $\gamma \in \{0.25, 1, 4\}$, and $r_L = r_S = r \in \{0.25, 1\}$. The remaining coefficients are $a_L = a_S = 0.05$, $c_L = c_S = 0.5$, $d_L = d_S = 2$, and $\kappa_L = \kappa_S = q_{\max} = 1$. This deterministic full-factorial design contains 96 synthetic coefficient sets. Every feasible interior, edge, and corner equilibrium is enumerated, including all real roots of the interior quadratic when $D < 0$; no sampled set has $D = 0$. Local stability uses the full Jacobian, with its general boundary form at edges and corners.

Each set is integrated from $(0.1, 0.9, 0.8)$ and $(0.9, 0.1, 0.2)$, giving 192 DOP853 trajectories with rtol $= 10^{-10}$, atol $= 10^{-12}$, maximum step 1, $T = 1000$, and output interval 0.5. An outcome is classified as converged only when the full-state maximum-norm error to an enumerated equilibrium is at most 10^{-6} at every output in the final 100 time units. The sampled T_ε retains the last-exit definition in Eq (4.1). Edge target equilibria have exactly one policy share on its boundary; corner target equilibria have both shares on their boundaries.

Table 12 separates feasibility from the sufficient certificate. The 36 feasible sets with $\Delta > 0$ give 72 interior outcomes, and a further 12 sets with $\Delta < 0$ give 24 locally stable interior outcomes. The six strongly competitive sets with $b_L = b_S = -3$ have unstable interior equilibria and two locally stable edge equilibria; the two starts select different edges in every set. Among the 42 sets without a feasible interior, 36 satisfy $\Delta > 0$ and yield 72 edge outcomes, whereas the remaining 6 yield 12 corner outcomes at $\left(1, 1, \gamma/(\gamma + 2r)\right)$. Thus, $\Delta > 0$ does not replace the feasibility check, and its failure does not imply local instability.

Table 12. Joint parameter sampling: equilibrium classification and trajectory outcomes.

Feasible interior / local status	Δ	Sets	Interior	Edge	Corner
Yes / stable	> 0	36	72	0	0
Yes / stable	< 0	12	24	0	0
Yes / unstable	< 0	6	0	12	0
No / not applicable	> 0	36	0	72	0
No / not applicable	< 0	6	0	0	12

Note: The last three columns count trajectories, with two initial states per coefficient set. Stable means that the equilibrium Jacobian is Hurwitz. No trajectory remains unresolved under the stated finite-horizon criterion. No sampled Δ equals zero.

Across the 192 trajectories, sampled T_ε ranges from 6.5 to 187.0. The maximum sampled bound violation is 9.992×10^{-15} without clipping, and the maximum equilibrium residual is below 1.125×10^{-15}. Radau checks of one stable interior case with $\Delta < 0$ and one competitive edge

case, using the same tolerances, horizon, and output interval, differ from DOP853 by at most 5.423×10^{-11}. This sample complements Figure 6 by coupling equilibrium classification to trajectories at every sampled parameter set. These finite-design counts do not estimate basin volumes or establish convergence for arbitrary parameters or initial states.

4.10. Numerical consistency and scope

The original suite has maximum equilibrium residual 4.441×10^{-16} and maximum Lyapunov-identity discrepancy 3.553×10^{-15} across applicable cases. The first 130 added DOP853 trajectories are finite and remain within the sampled bounds; selected equilibrium residuals are at most 3.886×10^{-16}. The extended unprojected LSODA trajectories have only numerical-scale departures: a maximum quality overshoot of 2.330×10^{-11} and policy undershoot of 3.426×10^{-17}. These values are retained rather than clipped.

The expanded study tests the gap between sufficient global stability and local stability, the role of asymmetric perturbations, nonlinear information response, and critical slowing. Grid coverage and finite trajectories do not establish a complete global bifurcation diagram. The empirical and structural limits of the model are consolidated in Section 5.3.

5. Conclusions and future research

This study develops a traffic-motivated evolutionary model coupling adaptive-policy shares in two CAV roles to information quality. Aggregate relative payoffs connect role interactions, information benefits, and adoption costs, while the quality state responds to recovery and policy-induced load. The analysis identifies conditions for coexistence and policy invasion and explains how these mechanisms shape long-run states and transients.

5.1. Main findings

First, a feasible interior equilibrium is globally asymptotically stable under the explicit dissipation condition, which also accommodates signed cross-role interactions. A separate Routh–Hurwitz criterion characterizes Hurwitz stability of the interior linearization. The numerical grid includes 88 locally stable cases outside the global sufficient region, while strong competitive effects produce an unstable interior point and distinct locally stable boundary equilibria. The 96-set joint-parameter sample further separates interior feasibility from convergence to stable boundary equilibria and confirms initial-state-dependent selection between competing edge equilibria.

Second, controlled payoff asymmetry reverses the ordering of role-specific adoption shares. Recovery-rate sweeps in three backgrounds increase equilibrium quality, whereas response-speed sweeps preserve equilibrium locations and can produce nonmonotonic convergence times. This separates equilibrium effects from transient effects.

Third, a common smooth strictly increasing quality-benefit function preserves the Lyapunov cancellation after its information term is modified. Normalized Hill responses change adoption levels and transient behavior, including initial policy disadvantage at low quality, without invalidating the sufficient stability result under its stated assumptions.

Fourth, policy invasion thresholds are roots of role-specific boundary payoff indicators. Their numerical values depend on the prescribed penetration-to-coefficient mapping: the two illustrative intercepts shift the policy invasion threshold from 0.5 to 0.4. Local invasion does not itself imply complete replacement of the fixed policy or convergence from arbitrary initial states.

Finally, near the original symmetric policy invasion threshold, the local spectral relaxation scale is proportional to $|p - 0.5|^{-1}$, whereas the critical symmetric trajectory decays algebraically. Extended integration confirms that failure to reach tolerance by $T = 2000$ is compatible with eventual convergence. Finite-observation convergence times must be distinguished from inverse eigenvalue scales.

5.2. Theoretical contributions and engineering implications

The relative-entropy construction cancels the policy–quality cross terms, yielding an explicit sufficient condition for global convergence. The Jacobian criterion separately identifies local stability outside that sufficient region. Signed interactions and common monotone information benefits reveal how the certificate depends on the coupling structure. Uniform scaling of the quality equation is a supporting corollary: It preserves equilibria and the certificate while changing the transient dynamics.

The model suggests examining information recovery jointly with the load generated by policy adoption and allowing payoff sensitivities to differ across traffic roles. Near a policy invasion threshold, both the relaxation rate and the finite observation horizon matter when interpreting numerical endpoints. The proposed coefficient-estimation route provides a basis for testing these mechanisms in microscopic traffic and communication simulations.

5.3. Limitations and future directions

The coefficients and information state have not been calibrated to road or communication data. The affine payoffs summarize utility differences rather than deriving them from a specific intersection game, and the scalar quality state does not resolve distinct sensing, communication, and computation impairments. Future calibration should use prespecified utility weights, independently varied operating conditions, and held-out traffic scenarios to assess identifiability and predictive validity. The numerical parameter ranges are illustrative, so their coverage is not a probability statement about real traffic. In particular, changing how payoff coefficients and unloaded quality depend on market penetration can shift a policy invasion threshold, remove it, or create several sign changes. These mathematical thresholds are not empirically established CAV deployment requirements.

The deterministic mean-field assumption omits finite-fleet fluctuations and spatially restricted interactions. Motivated by stochastic invasion studies [26,27], a separately specified imitation process could test policy-fixation probabilities and extinction times. Such effects cannot be inferred from repeated integrations of the same deterministic equations.

Extensions with communication delays, dynamic networks, heterogeneous recognition errors, or role-specific discontinuous information benefits require new stability analysis. Additional traffic roles and explicit trajectory-level objectives could connect policy evolution to lane changes, merging, and multi-vehicle intersections. Independent measures of traffic efficiency, safety, and fairness are required before assigning those engineering benefits to a dynamically stable policy distribution.

Use of Generative-AI tools declaration

The author used OpenAI's ChatGPT only for language polishing, grammatical checking, and improving the clarity of expression. All scientific content, theoretical derivations, results, analyses, and conclusions were independently completed, checked, and verified by the author.

Acknowledgments

The author would like to thank the National Science and Technology Council, Taiwan, for financially supporting this research (grant NSTC 115-2221-E-018-009).

Conflict of interest

The author declares no conflict of interest.

References

1. R. Li, Y. Liu, J. Sun, Y. Tian, Cooperative or competitive? Resolving social dilemmas in autonomous vehicles through evolutionary game theory, *Accid. Anal. Prev.*, **228** (2026), 108402. https://doi.org/10.1016/j.aap.2026.108402
2. S. -L. Lin, Y. -H. Chen, Deep learning–aided frequency-modulated continuous-wave radar for around-the-corner non-line-of-sight perception at urban intersections, *Comput. Model. Eng. Sci.*, **147** (2026), 37. https://doi.org/10.32604/cmes.2026.078862
3. S. -L. Lin, W. Wan, Intelligent connected vehicles, *Vehicles*, **8** (2026), 70. https://doi.org/10.3390/vehicles8040070
4. S. -L. Lin, An edge-assisted Internet-of-Vehicles computing framework for fair tail-risk allocation in cooperative autonomous driving, *Comput. Mater. Contin.*, **88** (2026), 11. https://doi.org/10.32604/cmc.2026.084571
5. J. S. Weitz, C. Eksin, K. Paarporn, S. P. Brown, W. C. Ratcliff, An oscillating tragedy of the commons in replicator dynamics with game-environment feedback, *Proc. Natl. Acad. Sci. U.S.A.*, **113** (2016), E7518–E7525. https://doi.org/10.1073/pnas.1604096113
6. A. R. Tilman, J. B. Plotkin, E. Akçay, Evolutionary games with environmental feedbacks, *Nat. Commun.*, **11** (2020), 915. https://doi.org/10.1038/s41467-020-14531-6
7. H. Gong, H. Xiang, Y. Wang, H. Gao, X. Meng, Strategy evolution of a novel cooperative game model induced by reward feedback and a time delay, *AIMS Mathematics*, **9** (2024), 33161–33184. https://doi.org/10.3934/math.20241583
8. M. Perc, A. Szolnoki, Coevolutionary games—A mini review, *BioSystems*, **99** (2010), 109–125. https://doi.org/10.1016/j.biosystems.2009.10.003
9. M. Perc, J. J. Jordan, D. G. Rand, Z. Wang, S. Boccaletti, A. Szolnoki, Statistical physics of human cooperation, *Phys. Rep.*, **687** (2017), 1–51. https://doi.org/10.1016/j.physrep.2017.05.004
10. M. Jusup, P. Holme, K. Kanazawa, M. Takayasu, I. Romić, Z. Wang, et al., Social physics, *Phys. Rep.*, **948** (2022), 1–148. https://doi.org/10.1016/j.physrep.2021.10.005

11. L. Cheng, M. Li, C. Tan, P. Huang, M. Zhang, R. Sun, Computational game-theoretic models for adaptive urban energy systems: A comprehensive review of algorithms, strategies, and engineering applications, *Arch. Computat. Methods Eng.*, **33** (2026), 2037–2114. https://doi.org/10.1007/s11831-025-10364-y
12. K. Wang, L. Cheng, M. Yin, K. Zhang, R. Wang, M. Zhang, et al., Evolutionary game theory in energy storage systems: A systematic review of collaborative decision-making, operational strategies, and coordination mechanisms for renewable energy integration, *Sustainability*, **17** (2025), 7400. https://doi.org/10.3390/su17167400
13. W. Schwarting, A. Pierson, J. Alonso-Mora, S. Karaman, D. Rus, Social behavior for autonomous vehicles, *Proc. Natl. Acad. Sci. U.S.A.*, **116** (2019), 24972–24978. https://doi.org/10.1073/pnas.1820676116
14. Y. Liu, X. Zhao, Y. Tian, J. Sun, Sociality Probe: Game-theoretic inverse reinforcement learning for modeling and quantifying social patterns in driving interaction, *IEEE Trans. Intell. Transp. Syst.*, **25** (2024), 20841–20853. https://doi.org/10.1109/TITS.2024.3461162
15. S. Shivshankar, A. Jamalipour, An evolutionary game theory-based approach to cooperation in VANETs under different network conditions, *IEEE Trans. Veh. Technol.*, **64** (2015), 2015–2022. https://doi.org/10.1109/TVT.2014.2334655
16. L. Ma, S. Qu, J. Ren, X. Zhang, Mixed traffic flow of human-driven vehicles and connected autonomous vehicles: String stability and fundamental diagram, *Math. Biosci. Eng.*, **20** (2023), 2280–2295. https://doi.org/10.3934/mbe.2023107
17. F. -C. Chou, A. Keimer, A. M. Bayen, Stability of ring roads and string stability of car following models, *Math. Control Relat. Fields*, **14** (2024), 1752–1775. https://doi.org/10.3934/mcrf.2024062
18. S. -L. Lin, Nonlinear stability and critical penetration analysis of mixed autonomous traffic with anticipative optimal-velocity control, *AIMS Mathematics*, **11** (2026), 15302–15323. https://doi.org/10.3934/math.2026629
19. T. Liard, R. Stern, M. L. D. Monache, A PDE-ODE model for traffic control with autonomous vehicles, *Netw. Heterog. Media*, **18** (2023), 1190–1206. https://doi.org/10.3934/nhm.2023051
20. P. De Leenheer, M. Schuster, H. Smith, Strong cooperation or tragedy of the commons in the chemostat, *Math. Biosci. Eng.*, **16** (2019), 139–149. https://doi.org/10.3934/mbe.2019007
21. Y. Wang, X. Meng, Evolutionary game dynamics of cooperation in prisoner's dilemma with time delay, *Math. Biosci. Eng.*, **20** (2023), 5024–5042. https://doi.org/10.3934/mbe.2023233
22. J. Liu, W. Hu, Bifurcation analysis of an asymmetric three-player three-strategy game with logit dynamics and reverse-cyclic mutations, *AIMS Mathematics*, **10** (2025), 22206–22220. https://doi.org/10.3934/math.2025988
23. H. Ohtsuki, M. A. Nowak, The replicator equation on graphs, *J. Theor. Biol.*, **243** (2006), 86–97. https://doi.org/10.1016/j.jtbi.2006.06.004
24. D. Madeo, C. Mocenni, J. C. Moraes, J. P. Zubelli, The role of self-loops and link removal in evolutionary games on networks, *Math. Biosci. Eng.*, **16** (2019), 5287–5306. https://doi.org/10.3934/mbe.2019264
25. A. S. Bratus, V. P. Posvyanskii, A. S. Novozhilov, A note on the replicator equation with explicit space and global regulation, *Math. Biosci. Eng.*, **8** (2011), 659–676. https://doi.org/10.3934/mbe.2011.8.659

26. A. Traulsen, M. A. Nowak, J. M. Pacheco, Stochastic dynamics of invasion and fixation, *Phys. Rev. E*, **74** (2006), 011909. https://doi.org/10.1103/PhysRevE.74.011909
27. J. F. Fontanari, Cooperation in the face of crisis: effect of demographic noise in collective-risk social dilemmas, *Math. Biosci. Eng.*, **21** (2024), 7480–7500. https://doi.org/10.3934/mbe.2024329
28. M. H. Mohd, J. Park, The interplay of rock-paper-scissors competition and environments mediates species coexistence and intriguing dynamics, *Chaos Solitons Fractals*, **153** (2021), 111579. https://doi.org/10.1016/j.chaos.2021.111579
29. J. Park, B. Jang, Structural stability of coexistence in evolutionary dynamics of cyclic competition, *Appl. Math. Comput.*, **394** (2021), 125794. https://doi.org/10.1016/j.amc.2020.125794
30. C. Du, Y. Lu, H. Meng, J. Park, Evolution of cooperation on reinforcement-learning driven-adaptive networks, *Chaos*, **34** (2024), 041101. https://doi.org/10.1063/5.0201968
31. J. Choi, S. Lee, H. Kim, J. Park, The role of recognition error in the stability of green-beard genes, *Evol. Lett.*, **7** (2023), 157–167. https://doi.org/10.1093/evlett/qrad012
32. G. Habib, Predicting saddle-node bifurcations using transient dynamics: A model-free approach, *Nonlinear Dyn.*, **111** (2023), 20579–20596. https://doi.org/10.1007/s11071-023-08941-6



AIMS Press

© 2026 the Author(s), licensee AIMS Press. This is an open access article distributed under the terms of the Creative Commons Attribution License (https://creativecommons.org/licenses/by/4.0)