Research Square

Preprints are preliminary reports that have not undergone peer review.
They should not be considered conclusive, used to inform clinical practice,
or referenced by the media as validated information.

Analytical and Finite-Population Properties of Defector-Responsive Strategy Switching

Aarush Gunjal

aarushgunjal1@gmail.com

Research Article

Keywords: evolutionary game theory, repeated Prisoner's Dilemma, replicator-mutator dynamics, state-dependent switching, cooperation, Moran process

Posted Date: August 4th, 2026

DOI: https://doi.org/10.21203/rs.3.rs-10502433/v1

License: This work is licensed under a Creative Commons Attribution 4.0 International License.
Read Full License

Additional Declarations: No competing interests reported.

Version of Record: A version of this preprint was published at Dynamic Games and Applications on October 31st, 2026. See the published version at https://doi.org/10.1007/s13235-026-00722-1.

Analytical and Finite-Population Properties of Defector-Responsive Strategy Switching

Aarush Gunjal[1*]

[1*]Independent Researcher, Newark, Delaware, USA; ORCID: 0009-0004-4007-532X.

Corresponding author(s). E-mail(s): aarushgunjal1@gmail.com;

Abstract

Defector-responsive switching from unconditional cooperation to Tit-for-Tat has been proposed as a way to rescue cooperation in the repeated Prisoner's Dilemma. The response rule itself was presented by Pal and coauthors; here we study analytical and finite-population consequences that were not characterized there. We first compare constant switching with switching proportional to the frequency of unconditional defectors. When maintaining the responsive mechanism reduces the payoff of unconditional cooperators, we derive an exact break-even cost for general benefit and cost parameters and give the conditions under which the relevant boundary equilibria resist invasion by defectors. We then implement a finite-population Moran process in which continuous switching intensities are mapped to valid per-update probabilities and cooperation is measured for a fixed period after rescue. Constant switching rapidly creates the retaliatory strategy but continues to remove unconditional cooperators after defectors disappear. Responsive switching can rescue more slowly at low intensity, but it shuts off after rescue and therefore preserves unconditional cooperation. These results identify a quantitative cost constraint and a finite-population tradeoff for an existing feedback mechanism, while also clarifying the biological assumptions required to interpret strategy switching as a phenotypic response.

Keywords: evolutionary game theory, repeated Prisoner's Dilemma, replicator-mutator dynamics, state-dependent switching, cooperation, Moran process

1 Introduction

The Prisoner's Dilemma expresses a conflict between individual and collective benefit: defection is individually favored in a single encounter, although mutual cooperation yields the larger joint return. Repeated interaction can change the evolutionary outcome because strategies may condition present behavior on earlier actions (Axelrod and Hamilton 1981; Axelrod 1984; Nowak 2006). Tit-for-Tat (TFT), which cooperates initially and then copies its opponent's preceding action, can punish unconditional defection while continuing to cooperate with unconditional cooperators (Nowak and Sigmund 1992). Later work has shown that the repeated game supports a much broader space of conditional and memory-dependent strategies (Stewart and Plotkin 2013).

Evolutionary game dynamics provide a natural framework for asking whether such behavior can protect a population dominated by unconditional cooperation. Replicator and selection-mutation equations describe the corresponding infinite-population dynamics (Taylor and Jonker 1978; Hofbauer and Sigmund 1998; Hofbauer 1985; Hadeler 1981; Page and Nowak 2002). Mutation can change both the number and stability of equilibria, including in multi-strategy and noisy social dilemmas (Antal et al. 2009; Chen et al. 2024). Finite-population formulations add demographic stochasticity and distinguish invasion, extinction, and fixation events (Moran 1962; Nowak et al. 2004; Traulsen et al. 2006; Fudenberg and Imhof 2006). In repeated games specifically, mutation can reorganize cycles of cooperation and defection or promote diverse cooperative communities (Imhof et al. 2005; Tkadlec et al. 2023).

Pal et al. (2025) studied a rescue mechanism with three strategies: unconditional cooperation (ALLC), TFT, and unconditional defection (ALLD). In their principal model, switching from ALLC to TFT occurs at a constant rate. TFT can suppress an ALLD invasion, but constant switching also continues after the invasion has ended. Importantly for the present study, Pal et al. also considered switching proportional to ALLD abundance in their electronic supplementary material, Figure S8. The feedback rule used here is therefore not claimed as a new mechanism.

Our contribution is instead to characterize two consequences of that rule. First, we add a payoff cost for maintaining a defector-responsive switching capability. We derive the exact cost at which the cooperative boundary equilibrium matches its constant-switching counterpart for arbitrary benefit b and cooperation cost c, and we state the stability restrictions under which this comparison is meaningful. Second, we revise the finite-population test. A continuous switching intensity is converted to a probability rather than used directly as one, and "survival quality" is evaluated after a fixed post-rescue horizon rather than at an absorbing ALLC state that does not exist under constant one-way switching.

We use the phrase *strategy switching* rather than genetic mutation throughout. The equations could describe mutation, learning, or reversible phenotypic change, but the proposed dependence on current defector abundance requires sensing or information about population state. Treating it as a phenotypic response is consequently the least restrictive interpretation. State-dependent perturbations have a longer history in evolutionary models (Bergin and Lipman 1996); the questions addressed here concern this particular feedback within the three-strategy rescue model.

2 Model

2.1 Payoffs and deterministic dynamics

Let x, y, and z be the frequencies of ALLC, TFT, and ALLD, respectively, with $x + y + z = 1$. We follow the payoff convention of Pal et al. (2025):

	ALLC	TFT	ALLD
ALLC	$b - c$	$b - c$	$-c$
TFT	$b - c$	$(b - c)/2$	0
ALLD	b	0	0

. (1)

Writing $q = b - c > 0$, the expected payoffs are

$$f_C = q(x + y) - cz, \qquad f_T = q(x + y/2), \qquad f_D = bx, \tag{2}$$

and $\bar{f} = xf_C + yf_T + zf_D$. One-way switching from ALLC to TFT gives

$$\begin{aligned} \dot{x} &= x(f_C - \bar{f}) - \mu x, \\ \dot{y} &= y(f_T - \bar{f}) + \mu x, \\ \dot{z} &= z(f_D - \bar{f}). \end{aligned} \tag{3}$$

We compare a constant rule, $\mu = \mu_0$, with the defector-responsive rule $\mu = \mu_0 z$.

To represent the cost of maintaining the responsive capability, the payoff of ALLC is reduced by $k \geq 0$. Thus, f_C is replaced by $f_C - k$ in its replicator term and the population mean becomes $\bar{f}_k = \bar{f} - kx$. The switching flow remains $\mu_0 zx$. This cost represents readiness to sense and respond, not the energetic cost of an individual switching event.

2.2 Finite-population process

The stochastic model is a frequency-dependent Moran birth-death process (Moran 1962; Nowak et al. 2004). It contains $N = 100$ individuals and begins with $N - 1$ ALLC and one ALLD. Payoffs are calculated against the other $N - 1$ individuals. At each update, a parent is sampled with weight

$$F_i = \exp(w\pi_i), \qquad w = 1, \tag{4}$$

where π_i is the expected game payoff, and a uniformly sampled individual is replaced. Exponential mapping keeps fitness positive even when game payoffs are negative. If the parent is ALLC, its offspring switches to TFT with probability

$$p(\lambda) = 1 - \exp(-\lambda), \tag{5}$$

where $\lambda = \mu_0$ for constant switching and $\lambda = \mu_0 z$ for responsive switching. Equation (5) prevents intensities greater than one from being misused as probabilities.

A trial is rescued when ALLD first reaches extinction and fails when ALLD fixes. A pre-rescue trial is capped at 100,000 updates. Rescued trials then continue for $50N$ updates, allowing us to measure the ALLC frequency after rescue even though constant switching makes all-ALLC nonabsorbing. We use 1,000 independent trials for each value of μ_0 and report means with normal-approximation 95% confidence intervals. No trial in the reported parameter sweep was censored.

3 Analytical Results

The boundary $z = 0$ permits a direct comparison of the two rules. It is also where their difference is clearest: constant switching remains active, whereas defector-responsive switching stops.

Proposition 1 *Assume $b > c > 0$ and write $q = b - c$.*

1. *Under constant switching, the nontrivial ALLC–TFT boundary equilibrium is*

$$x_F^* = 1 - \sqrt{\frac{2\mu_0}{q}}, \qquad y_F^* = \sqrt{\frac{2\mu_0}{q}}, \tag{6}$$

   and exists for $0 \le \mu_0 \le q/2$. Its transverse ALLD invasion exponent is

$$g_F(\mu_0) = c - b\sqrt{\frac{2\mu_0}{q}} + \mu_0. \tag{7}$$

2. *Under defector-responsive switching with readiness cost k, the nontrivial boundary equilibrium is*

$$x_R^* = 1 - \frac{2k}{q}, \qquad y_R^* = \frac{2k}{q}, \tag{8}$$

   and exists for $0 \le k \le q/2$. It resists an infinitesimal ALLD invasion if

$$k > \frac{cq}{b+c}. \tag{9}$$

3. *Whenever the two boundary equilibria are the relevant stable outcomes, their ALLC frequencies are equal at*

$$k^*(\mu_0) = \sqrt{\frac{q\mu_0}{2}}. \tag{10}$$

   For $cq/(b+c) < k < k^*(\mu_0)$, the responsive equilibrium contains more ALLC than the constant-switching equilibrium.

Proof. On $z = 0$, let $y = 1 - x$. For constant switching, the ALLC equation reduces to

$$\dot{x} = x\left(\frac{q}{2}y^2 - \mu_0\right), \tag{11}$$

which gives (6). At that equilibrium the mean payoff is $q - \mu_0$, so the growth rate of rare ALLD is $b(1 - y_F^*) - (q - \mu_0)$, yielding (7).

For responsive switching, $\mu_0 z = 0$ on the boundary. Coexistence requires equality of the cost-adjusted ALLC and TFT payoffs: $q - k = q(1 - y/2)$. This gives (8). The growth rate of rare ALLD is

$$b\left(1 - \frac{2k}{q}\right) - (q - k) = c - \frac{b + c}{q}k, \tag{12}$$

which is negative precisely under (9). Finally, equating (6) and (8) gives (10); the strict comparison follows from the same expressions. □

For the numerical values $b = 5$ and $c = 1$, (10) becomes $k^* = \sqrt{2\mu_0}$, while transverse stability of the responsive boundary requires $k > 2/3$. At the break-even value this condition holds when $\mu_0 > 2/9$. The constant boundary is transversely stable when $g_F(\mu_0) < 0$ and exists through $\mu_0 = 2$. These qualifications are essential: outside the joint boundary-stability regime, equating two boundary formulas does not compare the attracting states reached from the interior.

4 Numerical Results

4.1 Infinite-population comparison and cost check

We integrated (3) from $(x, y, z) = (1/3, 1/3, 1/3)$ to $t = 800$ with an eighth-order adaptive solver, relative tolerance 2×10^{-11}, and absolute tolerance 2×10^{-13}. The numerical implementation neither clips negative values nor renormalizes the state after integration; every terminal state was checked against the simplex constraint.

Figure 1A shows that constant switching has a narrow favorable region and then progressively removes ALLC as μ_0 increases. Responsive switching has a low-intensity tradeoff but rises to a high-ALLC equilibrium over the plotted range. For example, at $\mu_0 = 0.5$, the terminal ALLC frequencies are 0.500 under constant switching and 0.643 under responsive switching; at $\mu_0 = 1$, they are 0.293 and 0.701, respectively. At $\mu_0 = 2$, the constant-switching boundary reaches zero ALLC, while the responsive trajectory has terminal ALLC frequency 0.738.

For the costly model, we evaluated (10) at 31 values of μ_0 from 0.45 to 1.95. Figure 1B displays the predicted curve. The maximum absolute difference between the numerically integrated costly-responsive ALLC equilibrium and its constant-switching counterpart was 2.28×10^{-11}. This numerical check supports the implementation, while the result itself follows analytically from the proposition.

4.2 Finite-population tradeoff

Figure 2 separates three quantities that were previously conflated: whether ALLD becomes extinct, the composition at that first extinction time, and composition after a fixed post-rescue interval. At low intensity, constant switching rescues more reliably. For example, at $\mu_0 = 0.2$, the estimated rescue probabilities are 0.999 for constant

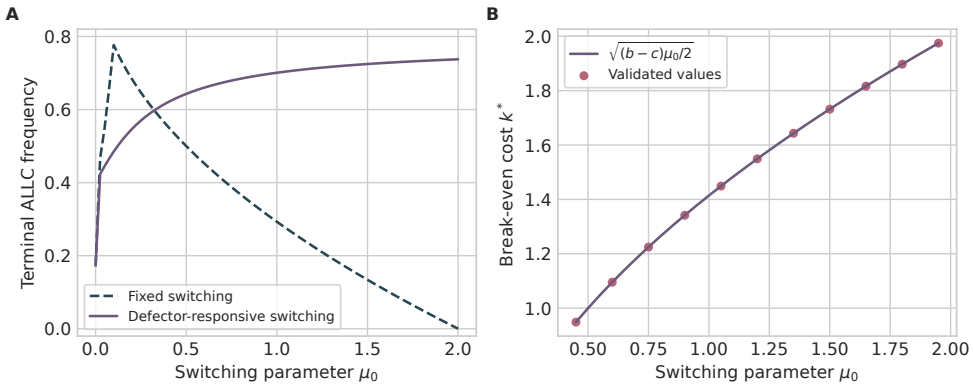


Fig. 1 Deterministic comparison for $b = 5$ and $c = 1$. (A) Terminal ALLC frequency from an interior initial condition under constant and defector-responsive switching. (B) Exact break-even readiness cost from (10); markers show parameters checked by numerical integration

switching and 0.749 ± 0.027 for responsive switching (95% confidence interval half-width). The responsive rule reacts weakly while the single invader has frequency $1/N$, explaining this delay.

The comparison reverses for post-rescue composition. At $\mu_0 = 0.6$, constant and responsive rescue probabilities are 1.000 and 0.988 ± 0.007, but their mean ALLC frequencies after $50N$ further updates are 0.370 ± 0.003 and 0.985 ± 0.008. At $\mu_0 = 1.4$, both rescue probabilities are one to the displayed precision, whereas post-rescue ALLC frequencies are 0.029 ± 0.003 for constant switching and 1.000 to three decimal places for responsive switching. Constant switching therefore buys rapid production of TFT at the cost of continuing conversion after ALLD is gone. The responsive rule increasingly combines high rescue probability with preservation of ALLC as its intensity rises.

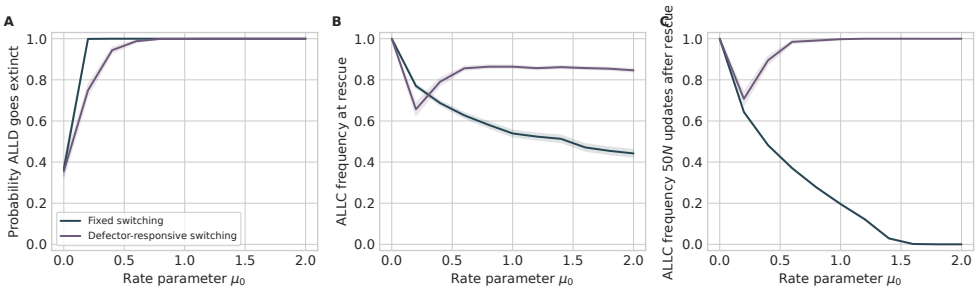


Fig. 2 Finite-population results for $N = 100$ and 1,000 trials per point. Lines show sample means and shading shows 95% confidence intervals. (A) Probability that ALLD reaches extinction before fixation. (B) ALLC frequency when ALLD first reaches extinction. (C) ALLC frequency after $50N$ additional updates

5 Discussion

The defector-proportional response is an appealing negative-feedback mechanism, but it is not new to this paper: Pal et al. (2025) presented it in their Supplementary Figure S8. The new analytical result here is the general break-even expression (10), together with the existence and invasion conditions needed to interpret it. The threshold makes the mechanism's cost constraint explicit. It also shows that the special-case expression $\sqrt{2\mu_0}$ is a consequence of choosing $b = 5$ and $c = 1$, rather than a parameter-free law.

The finite-population results refine the mechanism's interpretation. Constant switching can be more reliable when a defector is initially rare because it does not wait for a strong population-level signal. Responsive switching, however, shuts off exactly when $z = 0$ and consequently avoids continued conversion of ALLC after rescue. Reporting both rescue probability and post-rescue composition exposes this tradeoff. It would be misleading to report "ALLC fixation" under constant one-way switching, because all-ALLC is not absorbing: the next ALLC offspring can switch to TFT.

Several limitations restrict biological interpretation. The payoff matrix is a deliberately reduced representation of noisy repeated play, the switching direction is one-way, and the response uses instantaneous global knowledge of z. The readiness cost k is an abstract payoff penalty rather than a calibrated energetic measurement. The finite model also assumes that payoff enters fitness exponentially and that the same rate-to-probability mapping applies at every update. Thus, the results establish properties of a mathematical feedback model, not evidence for a particular molecular, microbial, or behavioral sensing mechanism. A biologically specific application would need local or noisy detection, a mechanistic response timescale, and parameters tied to data.

6 Conclusion

For the three-strategy rescue model, defector-responsive switching trades slower activation at low invader abundance for much better preservation of unconditional cooperation after rescue. In the deterministic costly model, the exact break-even readiness cost is the square root of one half the product of net cooperative benefit and constant switching intensity, subject to explicit boundary-stability conditions. Correctly converting rates to probabilities and measuring a finite post-rescue horizon makes the same tradeoff visible in finite populations. These results narrow the claim from proposing a mechanism to rigorously characterizing when an existing mechanism is advantageous.

Acknowledgements

The author thanks Dr. Saptarshi Pal for expert advice and helpful discussions concerning the evolutionary-game model.

Statements and Declarations

Funding

No funds, grants, or other support were received.

Competing Interests

The author has no relevant financial or non-financial interests to disclose.

Author Contributions

Aarush Gunjal performed the conceptualization, methodology, formal analysis, software development, numerical investigation, visualization, and manuscript preparation. The author read and approved the final manuscript.

Data Availability

The numerical data supporting the findings of this study are available in the public repository at https://github.com/aarushgunjal/defector_responsive_strategy_switching.

Code Availability

The Python source and figure-generation materials needed to reproduce the reported results are available at https://github.com/aarushgunjal/defector_responsive_strategy_switching. Running `python revised_analysis.py` regenerates the deterministic and finite-population outputs.

Use of Generative AI

Generative AI tools were used for language and presentation assistance during manuscript preparation. The author independently reviewed and verified the mathematical analysis, code, citations, numerical results, and final text and accepts full responsibility for the manuscript.

References

Axelrod, R., Hamilton, W.D.: The evolution of cooperation. Science **211**, 1390–1396 (1981) https://doi.org/10.1126/science.7466396

Antal, T., Traulsen, A., Ohtsuki, H., Tarnita, C.E., Nowak, M.A.: Mutation-selection equilibrium in games with multiple strategies. Journal of Theoretical Biology **258**, 614–622 (2009) https://doi.org/10.1016/j.jtbi.2009.02.010

Axelrod, R.: The Evolution of Cooperation. Basic Books, New York (1984)

Bergin, J., Lipman, B.L.: Evolution with state-dependent mutations. Econometrica **64**, 943–956 (1996) https://doi.org/10.2307/2171851

Chen, L., Deng, C., Duong, M.H., Han, T.A.: On the number of equilibria of the replicator-mutator dynamics for noisy social dilemmas. Chaos, Solitons & Fractals **180**, 114565 (2024) https://doi.org/10.1016/j.chaos.2024.114565

Fudenberg, D., Imhof, L.A.: Imitation processes with small mutations. Journal of Economic Theory **131**, 251–262 (2006) https://doi.org/10.1016/j.jet.2005.04.006

Hadeler, K.P.: Stable polymorphisms in a selection mutation model. SIAM Journal on Applied Mathematics **41**, 1–7 (1981) https://doi.org/10.1137/0141001

Hofbauer, J.: The selection mutation equation. Journal of Mathematical Biology **23**, 41–53 (1985) https://doi.org/10.1007/BF00276557

Hofbauer, J., Sigmund, K.: Evolutionary Games and Population Dynamics. Cambridge University Press, Cambridge (1998). https://doi.org/10.1017/CBO9781139173179

Imhof, L.A., Fudenberg, D., Nowak, M.A.: Evolutionary cycles of cooperation and defection. Proceedings of the National Academy of Sciences **102**, 10797–10800 (2005) https://doi.org/10.1073/pnas.0502589102

Moran, P.A.P.: The Statistical Processes of Evolutionary Theory. Clarendon Press, Oxford (1962)

Nowak, M.A.: Five rules for the evolution of cooperation. Science **314**, 1560–1563 (2006) https://doi.org/10.1126/science.1133755

Nowak, M.A., Sigmund, K.: Tit for tat in heterogeneous populations. Nature **355**, 250–253 (1992) https://doi.org/10.1038/355250a0

Nowak, M.A., Sasaki, A., Taylor, C., Fudenberg, D.: Emergence of cooperation and evolutionary stability in finite populations. Nature **428**, 646–650 (2004) https://doi.org/10.1038/nature02414

Pal, S., Lambert, A., Nowak, M.A.: Stabilizing unconditional cooperation. Proceedings of the Royal Society A **481**, 20240945 (2025) https://doi.org/10.1098/rspa.2024.0945

Page, K.M., Nowak, M.A.: Unifying evolutionary dynamics. Journal of Theoretical Biology **219**, 93–98 (2002) https://doi.org/10.1006/jtbi.2002.3066

Stewart, A.J., Plotkin, J.B.: From extortion to generosity, evolution in the iterated prisoner's dilemma. Proceedings of the National Academy of Sciences **110**, 15348–15353 (2013) https://doi.org/10.1073/pnas.1306246110

Tkadlec, J., Hilbe, C., Nowak, M.A.: Mutation enhances cooperation in direct reciprocity. Proceedings of the National Academy of Sciences **120**, 2221080120 (2023) https://doi.org/10.1073/pnas.2221080120

Taylor, P.D., Jonker, L.B.: Evolutionarily stable strategies and game dynamics. Mathematical Biosciences **40**, 145–156 (1978) https://doi.org/10.1016/0025-5564(78)90077-9

Traulsen, A., Nowak, M.A., Pacheco, J.M.: Stochastic dynamics of invasion and fixation. Physical Review E **74**, 011909 (2006) https://doi.org/10.1103/PhysRevE.74.011909