IJHSR
ijhighschoolresearch.org
RESEARCH ARTICLE

A Numerical Analysis of Stability in Zero-Sum Rock-Paper-Scissors Stochastic Replicator Dynamics

Kyle Yang

Deerfield Academy, 7 Boyden Lane, Deerfield, MA, 01342, USA; kyangsh2016@gmail.com

ABSTRACT: Stochastic differential equations have a wide range of applications across various disciplines, including finance, physics, and biology. There are usually two primary ways to evaluate them: Stratonovich, which samples midpoints, and Ito, which samples left endpoints. However, unlike deterministic calculus, these two schemes give different results upon evaluation. This article primarily analyzes the dynamics of the neutral equilibria of zero-sum replicator dynamics when subject to noise under both the Ito and Stratonovich schemes. We add noise because in real-world modeling scenarios, almost all natural systems are subject to some form of it. We find that intrinsically symmetric Brownian noise acts asymmetrically in zero-sum games, thereby decreasing a system's stability. Hence, neutral equilibria with noise behave like deterministic, asymptotically unstable equilibria. Additionally, the Stratonovich scheme exhibits greater variation than the Ito scheme in both mean and variance for all of the cases we study.

KEYWORDS: Chemistry, Organic Chemistry, Spectrophotometry, Carbinolamide, Kinetics.

Introduction

Stochastic Differential Equations:

This subsection largely follows Øksendal's textbook on SDEs, including notation and derivations.[1] In the modeling of various phenomena with random fluctuations, a special type of differential equation called a stochastic differential equation (SDE) is used. SDEs typically take the form:

$$dX(t) = \mu(t, X(t))\,dt \; + \; \sigma(t, X(t))\,d\xi(t) \tag{1}$$

Here, $X(t)$ is a variable probability distribution based on time, μ is a function from time and the distribution X (the state space) to codomain $\mathbb{R}$, σ is also a function from time and the state space to codomain $\mathbb{R}$, and $\xi(t)$ represents the random "noise" that is introduced to the system at time t. Essentially, the first term represents the driving movement of the system, while the second represents the random noise added.

The $\xi(t)$ can represent many types of noise. However, for this paper, we will assume $\xi(t)$ is a Wiener process (which models Brownian noise)—that is, we will assume $\xi(t) = W(t)$ to have the following properties, in accordance with Øksendal:[1]

i. $W(0) = 0$ almost surely (a.s.)

ii. $W(t)$ has independent increments

iii. $W(t+u) - W(t)$ is normally distributed with mean 0 and variance u: $W(t+u) - W(t) \sim \mathcal{N}(0,u)$

iv. $W(t)$ has a.s. continuous paths

Properties i) and iii) allow us to maintain symmetry, properties ii) and iii) give us nice probabilistic properties, and property iv) maintains well-definedness. Intuitively, the Wiener process helps us model a normal distribution whose variance increases linearly with time. At any time, we effectively 'sample' a random point from the distribution.

Having defined the terms in the SDE, we would now like to evaluate it. However, due to problems brought up by the Wiener process's lack of smoothness, the evaluated result changes based on the point sampled within each interval of a partition. We will demonstrate this difference between the Ito and Stratonovich schemes, which sample from the left endpoints and midpoints, respectively. These two schemes have different physical assumptions. While the Ito scheme accounts for the "unsmooth" nature of the noise, resulting in an extra term in its version of the chain rule, the Stratonovich scheme assumes a more symmetrical structure, maintaining the traditional deterministic chain rule. Following calculus intuition, we would like to "integrate both sides" of Equation 1 to determine the distribution of X. Since dW doesn't behave like a standard differential, we need to consider a (not necessarily even-spaced) partition of the interval $[0, T]$. Denote the finite collection of times dividing this interval as:

$$0 = t_1 < t_2 < \ldots < t_n = T$$

Now, we can assess the following discretized differentials and summations:

$$\Delta t_k = t_{k+1} - t_k, \qquad \Delta W_k = W(t_{k+1}) - W(t_k) \tag{2}$$

Left-endpoint (Ito) scheme:

$$X(t_n) = X(0) + \sum_{k=0}^{n-1} \mu(t_k, X(t_k))\,\Delta t_k + \sum_{k=0}^{n-1} \sigma(t_k, X(t_k)) \tag{3}$$

Midpoint (Stratonovich) scheme:

$$X(t_n) = X(0) + \sum_{k=0}^{n-1} \mu\left(\tfrac{t_k+t_{k+1}}{2}, \tfrac{X(t_k)+X(t_{k+1})}{2}\right)\Delta t_k + \sum_{k=0}^{n-1} \sigma\left(\tfrac{t_k+t_{k+1}}{2}, \tfrac{X(t_k)+X(t_{k+1})}{2}\right)\Delta W_k \tag{4}$$

Note that taking the limit as the size of each interval in the partition approaches 0 gives us the concept of "stochastic integrals." However, it is not too important for our purposes in this paper, as we will be adopting a discrete-time approach. We now demonstrate the difference between Ito (Equation 3) and Stratonovich (Equation 4) by taking their difference, which Øksendal terms the "correction term":[1]

$$\frac{1}{2}\sum_{k=0}^{n-1} \Delta X_k \Delta W_k \tag{5}$$

Generalizing these two schemes, we can sample from any point between intervals — say $\tau \in [t_k, t_{k+1})$. We denote α as

DOI: 10.36838/IJHSR87.8

the relative distance from t_k compared to t_{k+1}. In fact, regardless of which set of points we choose for the summation, our result remains mathematically consistent. However, despite mathematical consistency, the result of the summation yields nontrivial differences as we have demonstrated.

This paper will mainly analyze the differences between the $\alpha = 0$ (Ito) and $\alpha = ½$ (Stratonovich) schemes in the stochastic version of replicator dynamics. To contrast them, we will write Stratonovich SDEs with the notation $\circ\, dW_t$ and Ito SDEs with just dW_t, as Øksendal does.

Because the left–midpoint formulation of the Ito integral only uses information from time t_k to determine the state at time t_{k+1}, it doesn't "see into the future" (independent increments). In other words, this means the Ito scheme ensures that the mean is bounded and the conditional expectation up to a time is precisely the value at that time (martingale property). And, as mentioned previously, Ito integrals do not preserve the usual chain rule since an extra quadratic term from the Taylor expansion of $f(dX(t))$ appears due to the asymmetry of sampling the left midpoint and innate quadratic variation of W, giving the famous Ito's lemma:

$$df(X(t)) = \left[\mu(t)X(t)\frac{\partial f}{\partial X} + \tfrac{1}{2}\sigma(t)^2X(t)^2\frac{\partial^2 f}{\partial X^2}\right]dt + \sigma(t)X(t)\frac{\partial f}{\partial X}\,dW(t) \qquad (6)$$

A more detailed derivation using the Taylor expansion can be found in Øksendal's textbook.[1]

On the other hand, the Stratonovich integral preserves the classical chain rule. Because of its midpoint sampling, an extra quadratic term is generated, which cancels the quadratic term from the Taylor expansion. However, for the same reason, the Stratonovich integral does not have the martingale property; the midpoint sampling, in a certain sense, "peeks into the future."

Replicator Dynamics:

In this section, the definitions and derivations for deterministic replicator dynamics largely follow from Taylor and Jonker's 1978 paper and Hofbauer and Sigmund's 1998 book.[2,3] Replicator dynamics is a model that predicts frequencies of a set of competing phenotypes given a payoff matrix, which holds all the fitness information about the population ("rules" for when each pair of phenotypes "compete"). Mathematically, we describe this using a system of differential equations, with phenotype i adhering to:

$$dx_i = x_i\left[f_i(\mathbf{x}) - \phi(\mathbf{x})\right]dt, \qquad \phi(\mathbf{x}) = \sum_{j=1}^{n} x_j f_j(\mathbf{x}) \qquad (7)$$

Here, $\boldsymbol{x}$ is a column vector containing the frequencies of each phenotype; $\boldsymbol{\phi}$ is a function of $\boldsymbol{x}$, denoting the average fitness across all phenotypes; f_i denotes the fitness of phenotype i; and x_i is the frequency of phenotype i. The fitness of phenotype i is usually calculated by taking the i-th entry of the vector given by $\boldsymbol{Ax}$, where A is the payoff matrix. Essentially, the model compares each phenotype's fitness to the average fitness and determines the subsequent relative frequency of each phenotype based on the comparison. Furthermore, the model preserves the normalization of the x_i's, guaranteeing that they sum to 1. This can be easily shown by considering a function of the sum of all the x_i's and proving the derivative is 0.

We will follow Hofbauer and Sigmund's definitions of equilibrium points as those that have each $dx_i = 0$, *stable equilibria* as points that return to themselves under small perturbations, and *unstable equilibria* as points that become further *unstable*. Otherwise, we will call them *neutral*.[3]

The classic example of replicator dynamics is the *rock-paper-scissors* game, where A beats B, B beats C, and C beats A. It occurs frequently in nature, typically in the form of competing traits within a species. We will mainly study the zero-sum version in the rest of this paper, which means the winner and loser of each interaction gain/lose the same. As an example, we consider the following payoff matrix:

$$\begin{pmatrix} 0 & -1 & 1 \\ 1 & 0 & -1 \\ -1 & 1 & 0 \end{pmatrix}$$

Each ij-entry indicates the comparative fitnesses of phenotypes i and j, and more specifically, the outcome for type i. For example, when type 1 is compared to type 2, it will lose. Meanwhile, when type 1 is compared to type 3, it will win. Notice that to preserve the zero-sum nature, we must have an antisymmetric payoff matrix.

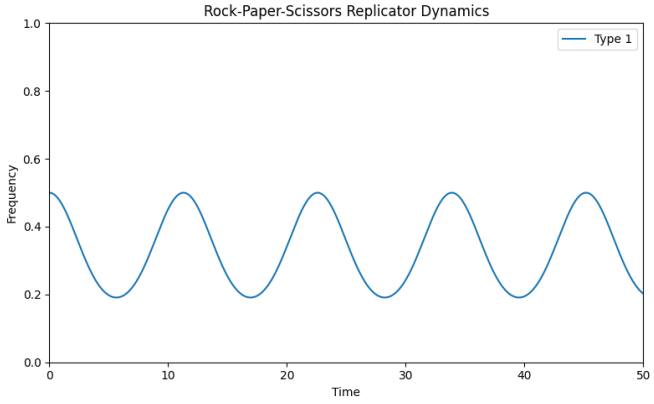

Figure 1: Frequency of phenotype x_1 as a function of time. The frequency never converges to the equilibrium point ⅓ because there is no attractive behavior in zero-sum replicator dynamics. Modeled using Heun's method, $\Delta t = 0.0001$.

Setting up the differential equations, initial values (0.5, 0.25, 0.25), and simulating the trajectories, we observe oscillatory behavior in Figure 1 due to the lack of attraction/repulsion exhibited by the neutral equilibrium point (⅓, ⅓, ⅓).

The majority of the following derivation of Ito replicator dynamics comes from Fudenberg and Harris's.[4] A more complete derivation with the individual calculations can be found in Delesse's work.[5] The derivation of the Stratonovich version is rather trivial, as it follows directly from the traditional chain rule.

When we consider adding a stochastic term to the differential equation, we can't directly add it, as the replicator differential equation is of frequency, not population. Thus, we must first consider the total populations and then convert them into their respective replicator equations. We have:

$$dr_i(t) = r_i(t)\left(\left[B(t,r(t)) - D(t,r(t))\right]dt + u_i[x(t)]dt + \sigma_i\,dW_i(t)\right) \qquad (8)$$

Where B is the background fitness ("base" number of successors), D is the death rate, u is the payoff function, σ_i is a scalar, and W_i is a Wiener process. Essentially, we are adding an extra stochastic term to each SDE in the system. Note that the quantity B-D differs from fitness, as we are considering the full population, not just frequencies. Now we would like to derive the differential equation for $x_i(t)$. Since we know $x_i(t) = r_i(t)/R(t)$, where $R(t)$ is the sum of all the ri's, we can apply the chain rule for Stratonovich and Ito's Lemma (Equation 6) for Ito on the function $f_i(r(t)) = r_i(t)/R(t)$, where $r(t)$ is a vector function containing the ri's. After expanding and simplifying:

Ito:

$$dx_i = x_i\left(f_i - \phi - \sigma_i^2 x_i + \sum_j \sigma_j^2 x_j^2\right)dt + x_i\left(\sigma_i\, dW_i - \sum_j \sigma_j x_j\, dW_j\right) \quad (9)$$

Stratonovich:

$$dx_i = x_i(f_i - \phi)\, dt + x_i\left(\sigma_i \circ dW_i - \sum_j x_j \sigma_j \circ dW_j\right) \quad (10)$$

Notice that the Ito version of replicator dynamics generates two extra terms in the drift term, which can be attributed to Ito's Lemma.

It has been shown that, given a stable equilibrium point, stochastic replicator dynamics frequently display behavior that also converges towards a neighborhood within that stable strategy. In other words, given that the noise is small enough, a stable equilibrium will still behave as a stable equilibrium. Additionally, a quantitative bound for the maximal magnitude of the noise that still maintains this stable behavior is known.6

It has also been demonstrated that, in certain instances, the stable states of the stochastic variant may differ substantially from those of the deterministic variant.[7] In other words, stochastic terms can cause the equilibria to shift. In addition, when there are multiple equilibrium points, it is also possible for trajectories to go back and forth between them, provided the noise and time are large enough.[8] The noise can push the trajectory into the attractive region of another equilibrium.

Together, these results demonstrate that equilibria cannot be treated in the same way when an extra stochastic term is added.

Despite the abundance of theory on the topic, very little numerical research has been done to observe the system's actualized behavior, especially in zero-sum stochastic models. To make new observations and test past claims regarding this topic, we will numerically simulate trajectories of zero-sum stochastic replicator dynamics, calculating the means and variances of phenotypes as a function of time. In the rest of this paper, we will analyze the speed of "divergence" for various magnitudes of noise in the stochastic version of the zero-sum rock-paper-scissors game when compared to the deterministic cycles illustrated in Figure 1.

Methods

To simulate individual trajectories for both Ito and Stratonovich processes, we used Heun's method with a time step of Δt = 0.001 in a Python script. We used Heun's method because it applies to both Ito and Stratonovich SDEs and maintains a relatively high level of accuracy.[9] Heun's method is as follows, in which each subsequent step takes the average of a "predictor" and "corrector":

$$\tilde{X}_{n+1} = X_n + f(t_n, X_n)\,\Delta t + G(t_n, X_n)\,\Delta W_n \quad (11)$$

$$X_{n+1} = X_n + \tfrac{1}{2}[f(t_n, X_n) + f(t_{n+1}, \tilde{X}_{n+1})]\Delta t + \tfrac{1}{2}[G(t_n, X_n) + G(t_{n+1}, \tilde{X}_{n+1})]\Delta W_n \quad (12)$$

We will assume that the variance (σ) of each phenotype is the same, ranging from the set: {0.05, 0.10, 0.15, 0.20, 0.25}. For each σ, we ran N = 10000 simulations for a duration of T = 50.0 with the initial condition $\boldsymbol{x}$ = (0.5, 0.25, 0.25). We ran this simulation several times for each σ and found similar results. Furthermore, the general trends observed in our N = 1000 test runs are almost identical to those in our N = 10000 runs. These suggest that our results are converging.

Each random variable W was generated using CuPy's built-in random number generator, which adheres to the rules specified above for a Wiener process. Since each path runs independently, using parallel computing can significantly speed up a simulation while retaining the simulation's integrity. Therefore, we utilized a GPU and CuPy's intrinsic GPU optimizations to accelerate the simulation.

To generate a frequency-time heatmap, we partitioned the time axis into 500 sections and the frequency axis into 50 sections, calculating probabilities by counting the number of trajectories passing through each partitioned section relative to the entire vertical strip. We collected data on the simulated mean and variance of the dataset in addition to the heatmaps.

Results and Discussion

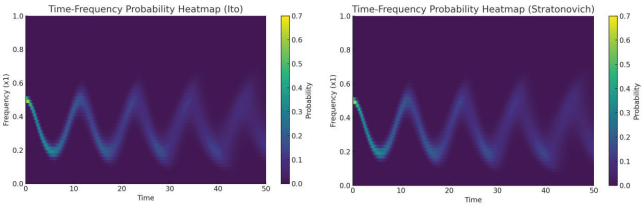

(A) (B)

Figure 2: Simulated results for σ = 0.05. The colors in the heatmap also begin to fade over time, indicating dispersion in individual trajectories. However, the general neighborhood of trajectories is still visible. A) Ito heatmap B) Stratonovich heatmap.

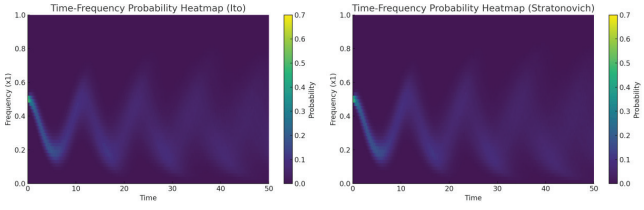

(A) (B)

Figure 3: Simulated results for σ = 0.10. The colors in the heatmap also begin to fade over time, indicating dispersion in individual trajectories. The general neighborhood of trajectories is still slightly visible. A) Ito heatmap B) Stratonovich heatmap.

DOI: 10.36838/IJHSR87.8

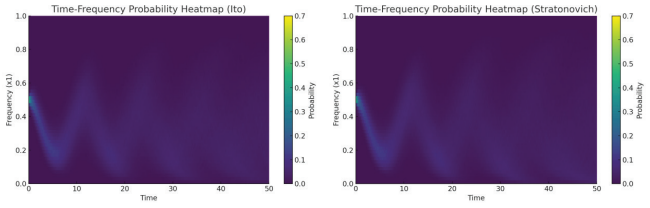

(A) (B)

Figure 4: Simulated results for σ = 0.15. The colors in the heatmap also begin to fade over time, indicating dispersion in individual trajectories. The general neighborhood of trajectories becomes unclear at about halfway, but generally concentrates more at lower frequencies. A) Ito heatmap B) Stratonovich heatmap.

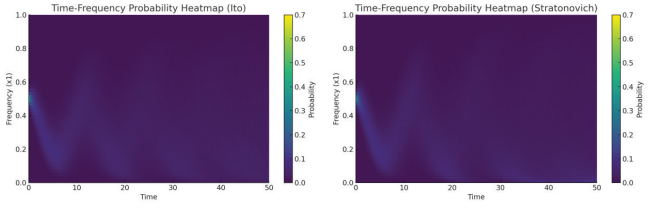

(A) (B)

Figure 5: Simulated results for σ = 0.20. The colors in the heatmap also begin to fade over time with a larger concentration in a low-frequency range, indicating dispersion in individual trajectories, which favors larger oscillations. The general neighborhood of trajectories becomes unclear quickly. A) Ito heatmap B) Stratonovich heatmap.

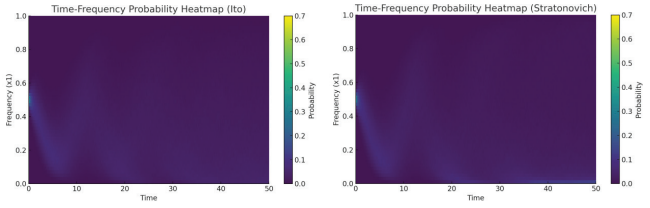

(A) (B)

Figure 6: Simulated results for σ = 0.25. The colors in the heatmap also begin to fade over time, with a larger concentration in a low-frequency range, indicating dispersion in individual trajectories, which favors larger oscillations. The general neighborhood of trajectories becomes unclear quickly. A) Ito heatmap B) Stratonovich heatmap.

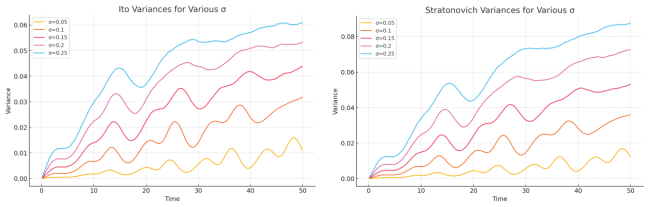

(A) (B)

Figure 7: Variances for various values of σ. They remain mostly evenly spaced in the time sampled. They all appear to dampen over time, although at different rates. A) Ito variances B) Stratonovich variances

Variance Behavior:

As seen in Figures 2-6 (a) and Figures 2-6 (b), over time, the paths disperse and stray from the orbits of the deterministic version. It is also clear that the size of the oscillations increases. However, the rate of this increase is variable. In Figures 7 (a)-(b), we observe that the rate at which variance increases over time is roughly proportional to the size of σ before beginning to dampen.

Notice that in Figure 7, the variances mostly remain evenly spaced and exhibit the behavior of "flattening out" as time goes on. We can attribute the second result to the fact that frequency is bounded between 0 and 1, and hence its variance is also bounded.

Furthermore, in Figure 7, oscillations are mirroring those in Figures 2-6 (a) and Figures 2-6 (b). This can be attributed to the fact that trajectories are "most misaligned" when the graph reaches its peak, as seen by the wider range of values in Figures 2-6 (a) and Figures 2-6 (b) at that point.

Ito vs. Stratonovich:

From observation, the Stratonovich variance is greater than the Ito variance. This result is due to the extra terms in the Ito version (especially the negative term), which act as an opposing force to any phenotype getting too big. We will provide a mathematical bound for when the Ito form's extra terms are guaranteed to limit the SDE's growth.

In the SDE, notice that the - $\sigma_i^2 x_i$ has a linear xi, while every other $\sigma_j^2 x_j^2$ has quadratic x_j. Also, note that each $0 < x_k < 1$, and they sum to 1. We would like the negative term to "overpower" the positive terms. Equivalently, we need that $\sigma_i^2 x_i > \Sigma\, \sigma_j^2 x_j^2$. Specifically, when we assume each $\sigma_i = \sigma$, as we did in our simulation, we can cancel the σ's as follows:

$$x_i > \sum_{j \neq i} x_j^2 \tag{13}$$

Since each $x_j > 0$, which is true by hypothesis, we can apply the ℓ1-ℓ2 norm inequality:

$$\left(\sum_{j \neq i} x_j\right)^2 \geq \sum_{j \neq i} x_j^2 \tag{14}$$

Since the sum of the x_i's is 1, we can make a substitution for the left-hand side of equation 14 and use the following inequality chain relating equations 13 and 14 to determine a bound for x_i :

$$x_i > (1 - x_i)^2 = \left(\sum_{j \neq i} x_j\right)^2 \geq \sum_{j \neq i} x_j^2 \tag{15}$$

Since any x_i satisfying the first inequality in equation 15 will automatically satisfy equation 13, we can solve this inequality to generate a bound for when the Ito terms are *guaranteed* to limit the SDE's growth. It is clear by algebra that this inequality is satisfied when $x_i > 0.382$. Interestingly, this is equal to one minus the inverse of the golden ratio. Thus, when some phenotype xi even slightly dominates the other phenotypes, the extra term in the Ito scheme dampens further growth. Additionally, we hypothesize that as time increases, the Ito scheme (Equation 9) will exhibit a lower maximum variance than the Stratonovich scheme (Equation 10).

Mean Behavior:

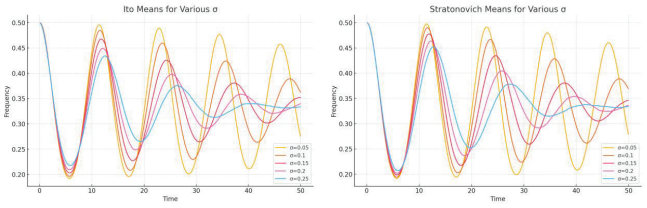

(A) (B)

Figure 7: Means trajectories for various values of σ. The rate of convergence increases as the magnitude of the noise increases. A) Ito means B) Stratonovich means.

While the variances and sizes of oscillations tend to increase for each σ, the means tend to stabilize towards the center equilibrium ⅓, as shown in Figure 8. Rather unintuitively, the speed of stabilization increases with σ, as shown in Figure 8; notice that as σ increases, the mean converges to ⅓ faster.

This effect can be attributed to the fact that noise can cause the periods of the cycles to change. This results in an average that involves the sum of many misaligned cycles, which cancel each other out and converge towards equilibrium. Hence, we have shown that a converging average does not necessarily indicate converging trajectories; in fact, a converging average often indicates increasingly chaotic trajectories. Essentially, the system exhibits some ergodicity.

Implications:

Together, these results indicate that asymptotically neutral equilibrium points are closer in nature to asymptotically unstable points. In other words, when a system has no "preferences," it tends to grow more chaotic rather than remain stable. The noise does not act symmetrically.

Furthermore, these results indicate that neutral equilibria and their associated cycles are susceptible to noise-induced changes, as the system lacks an intrinsic counterbalance. This can be observed in the dispersive behavior shown in Figures 2-6 and the steady increase in variance illustrated in Figure 7.

Since underlying neutral equilibrium points tend to be unstable, this result implies that cycles either have underlying stable behavior or have a near-negligible noise (small-noise limit). In the first case, the noise counteracts the underlying stability. In the second case, the system is inherently neutral in nature and typically adopts the deterministic variant.

Similar cyclic behavior can be commonly observed in nature's ecosystems, whether it be in the form of predator-prey cycles or competing phenotypes in a rock-paper-scissors-type relationship. As we have shown, these structures are incredibly brittle and easily disturbed by changes in the magnitude of noise to which the system is subject. In ecosystems, climate change, overfishing, and other human activities can all contribute to this environmental variance. Although noise does not immediately cause extinction events (i.e., one frequency approaches 0), increases in variance and cycle amplitude, seen in Figures 2-7, can significantly increase the risk. Notice that in each heatmap, a large percentage of the data concentrates at the more extreme frequencies. This behavior renders ecosystems far more susceptible to disasters, which can turn these extreme events into extinction-level occurrences.

Regarding the Ito vs. Stratonovich question, both the theory and results indicate that the Ito scheme is more resilient to noise than the Stratonovich scheme. These differences result from the assumed properties of the noise. While the Ito scheme assumes a martingale, the Stratonovich scheme assumes smoothness in the noise. Ultimately, the better model depends on the underlying conditions of the noise, which should be chosen carefully.

Further Questions:

Although we have provided a numerical approximation for the effect of noise on zero-sum games, several questions remain unanswered. What is the generalized behavior and stability of the system for any generic antisymmetric payoff matrix? What if the noise vector is asymmetric, where some $\sigma_i \neq \sigma_j$? Is there an explicit formula for the mean's speed of convergence and the variance's speed of increase? How do we express the sizes of the oscillations for the variance and mean graphs quantitatively? Which more accurately models noise in the context of rock-paper-scissors: Ito or Stratonovich?

■ Conclusion

We have illustrated that despite a converging mean, individual trajectories may diverge wildly within a system. In zero-sum games, adding a stochastic term often increases the system's instability, resulting in larger oscillations. In these cases, the magnitude of the system's variance also increases proportionally with the underlying variance before dampening. Specifically, our studied payoff matrix is

$$\begin{pmatrix} 0 & -1 & 1 \\ 1 & 0 & -1 \\ -1 & 1 & 0 \end{pmatrix}$$

Furthermore, as σ increases, the mean converges more quickly towards the neutral equilibrium point, with each peak decreasing steadily before evening out.

In both mean and variance, Ito dynamics (Equation 9) display less movement than Stratonovich dynamics (Equation 10), attributed to the "Ito correction" resulting from Ito's Lemma (Equation 6). This difference is non-negligible and should be carefully considered in modeling settings. We have also given an inequality, xi > 0.382, for the threshold at which the Ito scheme is guaranteed to limit growth more than the Stratonovich scheme.

Since symmetric Brownian noise acts asymmetrically on neutral equilibrium points and cycles, such equilibria behave very similarly to intrinsically asymptotically unstable equilibria. Hence, it is unlikely that observed neutral equilibria are intrinsically neutral unless the noise is sufficiently small. Additionally, without any attractive counterbalance, these equilibria are very susceptible to noise, which can significantly alter a trajectory's path. These alterations typically manifest as increased cycle amplitudes and concentrations at more extreme frequencies. In nature, environmental variance often comes from factors such as climate change, overfishing, and other human activities. These can significantly offset a system's balance and lead to extinction events.

■ Acknowledgments

The author thanks Professor J.S.W. Lamb for his guidance and mentorship.

■ References

1. Øksendal, B., Stochastic Differential Equations: An Introduction with Applications, 6th ed.; Springer: 2003.

DOI: 10.36838/IJHSR87.8

2. Taylor, P. D.; Jonker, L. B. Evolutionarily Stable Strategies and Game Dynamics. Mathematical Biosciences 1978, 40, 145–156.
3. Hofbauer, J.; Sigmund, K., Evolutionary Games and Population Dynamics; Cambridge University Press: Cambridge, 1998.
4. Fudenberg, D.; Harris, C. Evolutionary Dynamics with Aggregate Shocks. Journal of Economic Theory 1992, 57, 420–441.
5. Delesse, F. Stochastic Replicator Dynamics, Mémoire d'initiation à la recherche, PSL Research University – Université Paris-Dauphine, 2016.
6. Imhof, L. A. The Long-run Behavior of the Stochastic Replicator Dynamics. The Annals of Applied Probability 2005, 15, 1019–1045.
7. Foster, D. P.; Young, H. P. Stochastic Evolutionary Game Dynamics. Theoretical Population Biology 1990, 38, 219–232.
8. Avrachenkov, K.; Borkar, V. S. Metastability in Stochastic Replicator Dynamics. Dynamic Games and Applications 2019, 9, 366–390.
9. Hairer, E.; Nørsett, S. P.; Wanner, G., Solving Ordinary Differential Equations I: Nonstiff Problems, 2nd ed.; Springer: 1993.

Author

Kyle Yang attends Deerfield Academy (G11) in Massachusetts, USA. He is interested in both pure and applied mathematics, especially statistical theory, abstract algebra, and number theory. He plans on studying mathematics in college.