Minimax Theory and its Applications
Volume 12 (2027), No. 1, 24 to 28

Risk-aware minimax equilibrium computation for stochastic multi-agent systems using proximal game dynamics

Dr. Nidhi Mishra[1], Aakansha Soy[2]

[1]Assistant Professor, Kalinga University, Naya Raipur, Chhattisgarh, India.
ku.nidhimishra@kalingauniversity.ac.in, https://orcid.org/0009-0001-9755-7950
[2]Assistant Professor, Kalinga University, Naya Raipur, Chhattisgarh, India.
ku.aakanshasoy@kalingauniversity.ac.in, https://orcid.org/0009-0002-1955-6909

Abstract:
Agents in a stochastic game rarely consider only the mean. In a firm where a bad quarter can kill you, the tail of the loss distribution matters. When every agent minimises its own conditional value-at-risk, its problem reduces to a minimax problem against an adversarial reweighting of nature. That game has a unique Nash equilibrium which forms a system of coupled minimax problems. We formulate this risk-aware game as a stochastic variational inequality in an expanded space including each agent's value-at-risk level. We derive existence conditions, and study regularised game dynamics: an outer proximal-point iteration with inner regularised games solved by stochastic approximation. The regularisation both stalls the effect of the inner step, which is an outer stochastic gradient step, and caps the step size so that the regularised game is insensitive to the step size. In a four-firm Cournot market subject to correlated demand and cost shocks, the regularisation sustains a constant step size range under which the methods behave accurately and is robust to poorly tuned step sizes. For a commonly used initial step size, the regularised method's worst final distance to equilibrium outperformed the plain stochastic projected gradient method's (0.14 versus 0.66). At best-tuned step size, the plain method was more accurate (0.03 versus 0.06). The model reveals that risk aversion reduces total output by 37% and increases total expected profit by 34%, relative to risk-neutrality, making risk aversion an implicit policy against competition.

Keywords*:* stochastic Nash equilibrium; conditional value-at-risk; minimax; proximal point method; stochastic approximation; variational inequality

MSC 2020: 91A10, 91A15, 90C47, 49J40, 62L20

1. Introduction

Consider Suppose four generators bid quantities in a market with a stochastic demand shock and each has a stochastic fuel-cost shock that is high precisely when demand is low. In the classical model each firm maximizes the expected profit and the resulting solution is a Cournot-Nash equilibrium (Nash, 1950). Actual firms are not like this: credit covenants, margin calls and managers' incentives all penalise undesirable results and reward good ones. A more realistic model allows each firm to minimize some measure of its loss.
We rely on CVaR for two reasons. First, it is coherent (see Artzner et al. (1999)) so it has a dual representation as a worst case expectation. Minimising CVaR is thus a minimax problem where the adversary optimally reweights the scenario distribution but with a bounded density ratio. Second, Rockafellar and Uryasev (2000) proved that CVaR minimisation is a convex program in the decision and one other auxiliary scalar and thus manageable inside an equilibrium calculation.
This object, that we call risk-aware minimax equilibrium, leads to two issues of practical concern.
Q1. Does an equilibrium exist, and in what sense is it well behaved? Existence for convex stochastic Nash games is classical (Rosen, 1965; Ravat & Shanbhag, 2011). The CVaR transformation preserves convexity in one's own actions, but can wreck the smoothness that many solvers rely on.
Q2. How to estimate from samples only? Stochastic approximation for monotone variational inequalities is fairly mature (Koshal et al., 2013; Juditsky et al., 2011; Mertikopoulos & Zhou, 2019). The CVaR gradient is a predictor-weighted quantity whose variance is amplified by up to the inverse of 1-α; so approaches that work for a single α-value may be slow to make progress or diverge at others.
Our contributions. (i) We derive the equilibrium as a monotone-type stochastic variational inequality in the extended variables (qi, ti), with t_i being the agent i 's level-at-risk. (ii) We investigate a proximal game dynamic, a stochastic and inexact variant of Rockafellar's (1976) proximal point method. Its regularisation imposes an inner step size cap and we pinpoint this as the reason for the overall method's robustness. (iii) A numerical experiment on a correlated-shock Cournot market evaluates the economic impact of risk aversion as

well as the robustness-peak accuracy trade-off of the two solvers.

2. Model

2.1. Players, actions, losses

There are N agents. Agent i chooses $q_i \in K_i = [0, \bar{q}_i]$, and $q = (q_1, \dots, q_N)$. Uncertainty is a random vector ω on a probability space $(\Omega, \mathcal{F}, P)$. Agent i's random loss is $L_i(q_i, q_{-i}; \omega)$, which we assume is convex and continuously differentiable in q_i for every ω and integrable for every q.

2.2. Risk preference as a minimax problem

Fix a confidence level $\alpha \in [0,1)$. Agent i minimises

$$\rho_i(q) = \mathrm{CVaR}_\alpha[L_i(q;\omega)] = \sup_{Q \in \mathcal{Q}_\alpha} \mathbb{E}_Q[L_i(q;\omega)], \qquad \mathcal{Q}_\alpha = \left\{ Q \ll P:\ 0 \le \frac{dQ}{dP} \le \frac{1}{1-\alpha} \right\}. \qquad (1)$$

The second equality is the dual representation of a coherent risk measure (Artzner et al., 1999). Each agent therefore plays its own zero-sum game against a "nature" player who may inflate the probability of any event by a factor of at most $1/(1-\alpha)$. At $\alpha = 0$ nature has no power and ρ_i is the expectation. As $\alpha \to 1$ nature concentrates all mass on the worst outcomes.

Definition 1. A profile $q^* \in K = \prod_i K_i$ is a *risk-aware minimax equilibrium* if, for every i,

$$q_i^* \in \arg\min_{q_i \in K_i} \sup_{Q \in \mathcal{Q}_\alpha} \mathbb{E}_Q[L_i(q_i, q_{-i}^*; \omega)].$$

2.3. Extended formulation

The supremum in (1) is not convenient for computation. Rockafellar and Uryasev (2000) give the equivalent minimisation

$$\mathrm{CVaR}_\alpha[L] = \min_{t \in \mathbb{R}} \Phi(t, L), \qquad \Phi(t, L) = t + \frac{1}{1-\alpha}\, \mathbb{E}[(L-t)_+], \qquad (2)$$

whose minimisers include the value-at-risk of L. Substituting (2) gives an extended game. Player i chooses $x_i = (q_i, t_i) \in X_i = K_i \times [-T, T]$, with T large enough to contain every relevant quantile, and minimises $\Phi_i(x_i; q_{-i}) = \Phi\big(t_i, L_i(q;\omega)\big)$. The extended game and Definition 1 have the same equilibria in q, because the minimisation in t_i is separable and involves only agent i.

2.4. The stochastic pseudo-gradient

For a sample ω, a stochastic subgradient of Φ_i with respect to x_i is

$$G_i(x;\omega) = \begin{pmatrix} \dfrac{1}{1-\alpha}\, \mathbf{1}\{L_i > t_i\}\ \partial_{q_i} L_i(q;\omega) \\ 1 - \dfrac{1}{1-\alpha}\, \mathbf{1}\{L_i > t_i\} \end{pmatrix}. \qquad (3)$$

Stacking gives $G(x;\omega)$ and its mean $F(x) = \mathbb{E}[G(x;\omega)]$. Note the factor $1/(1-\alpha)$ in (3). At $\alpha = 0.9$ only about 10% of samples contribute, each multiplied by 10. The estimator is unbiased, but its variance grows roughly like $1/(1-\alpha)$. This is the computational cost of being risk-aware.

3. Existence and structure

Assumption A. (A1) Each K_i is a nonempty compact interval. (A2) For each ω, $L_i(\cdot, q_{-i}; \omega)$ is convex on K_i. (A3) $L_i(q;\omega)$ is continuous in q, and $\sup_{q \in K} |L_i(q;\omega)|$ is integrable.

Proposition 1 (existence). Under Assumption A, a risk-aware minimax equilibrium exists for every $\alpha \in [0,1)$.

Proof. CVaR is monotone and convex as a functional of the loss, and L_i is convex in q_i. Hence $\rho_i(\cdot, q_{-i})$ is convex. By (A3) and dominated convergence, $\mathbb{E}_Q[L_i]$ is continuous in q uniformly over $Q \in \mathcal{Q}_\alpha$, because the density ratio is bounded by $1/(1-\alpha)$. The supremum (1) is then continuous in q. Each player's cost is continuous and convex in its own action on a compact convex set, so Rosen's theorem (1965, Thm. 1), applied to the payoffs $-\rho_i$, gives an equilibrium. ▫

Proposition 2 (variational characterisation). Suppose in addition that $L_i(q;\omega)$ has a continuous distribution for each q. Then Φ_i is continuously differentiable in x_i, and x^* is an equilibrium of the extended game if and only if it solves the variational inequality

$$\text{find } x^* \in X: \qquad \langle F(x^*), x - x^* \rangle \ge 0 \quad \forall\, x \in X = \prod_i X_i\,. \qquad (4)$$

Proof. Differentiability of $\mathbb{E}[(L-t)_+]$ in t under a continuous distribution is standard, and differentiability in q_i follows from (A2)–(A3). Each player's problem is convex, so it is equivalent to its first-order condition. Stacking the N conditions gives (4), as in the general reduction of Facchinei and Kanzow (2007). For losses with atoms, F becomes a set-valued subdifferential map, and (4) holds in the generalised sense of Ravat and Shanbhag (2011). ▫

Monotonicity. Convergence of first-order methods depends on whether F is monotone, that is, $\langle F(x) - F(y), x - y \rangle \ge 0$. Two remarks are in order.

Remark 1 (risk-neutral case). If $\alpha = 0$ and the market is a linear Cournot market, $L_i = q_i(c_i - a + bQ - \xi) + 1/2\, \gamma_i \zeta q_i^2$, the t-block decouples. The q-block has the constant Jacobian

$b(I + \mathbf{1}\mathbf{1}^\top) + \mathrm{diag}(\gamma_i\, \mathbb{E}[\zeta])$, which is positive definite. F is then strongly monotone, the equilibrium is unique, and plain stochastic approximation converges (Koshal et al., 2013).

Remark 2 (risk-averse case). For $\alpha > 0$ the tail event $\{L_i > t_i\}$ depends on the rivals' quantities, and monotonicity of the extended map is no longer automatic. Rosen's diagonal strict convexity (1965) is a sufficient condition, but it is hard to check from samples. In this paper we take monotonicity of F on X as a standing hypothesis (Assumption M). In Section 5 we check the solution numerically: Gauss–Seidel best-response iteration converged to a Nikaido–Isoda gap below 10^{-12} at every α tested.

The proximal algorithm to be presented in the next section was specifically devised for this scenario. It requires monotonicity but not strong monotonicity, and it provides the required curvature by itself.

4. Proximal game dynamics

4.1. Two solvers

Stochastic projected gradient (SPG). This is the baseline. Every agent updates simultaneously using a mini-batch $\omega_k^1, \dots, \omega_k^B$:

$$x^{k+1} = \Pi_X\left(x^k - \frac{\gamma_0}{\sqrt{k}} \cdot \frac{1}{B} \sum_{b=1}^{B} G\left(x^k; \omega_k^b\right) \right).$$

Proximal game dynamics (PGD). An outer sequence of anchors z^m is updated by approximately evaluating the resolvent of $F + N_X$, where N_X is the normal cone of X:

$$z^{m+1} \approx J_\lambda(z^m), \qquad J_\lambda(z) = \text{the unique } x \in X \text{ with } 0 \in F(x) + N_X(x) + \frac{1}{\lambda}(x - z). \qquad (5)$$

In game terms, (5) is a modified game in which each agent pays an extra cost $\frac{1}{2\lambda} \| x_i - z_i^m \|^2$ for moving away from its previous anchor. Each agent stays rational and only becomes reluctant to move far in one round.

The inner game is solved by stochastic approximation:

$$x^{j+1} = \Pi_X\left(x^j - \gamma_j \left[\frac{1}{B} \sum_b G\left(x^j; \omega_j^b\right) + \frac{1}{\lambda}(x^j - z^m) \right] \right), \qquad \gamma_j = \min\{\gamma_0,\ \lambda/j\}, \qquad (6)$$

for $j = 1, \dots, n_m$. The new anchor z^{m+1} is the average of the second half of the inner iterates.

4.2. Why the regularisation helps

Two facts carry the argument.

Fact 1: the inner game is strongly monotone. If F is monotone (Assumption M), the inner map $F + \frac{1}{\lambda}(\cdot - z^m)$ is strongly monotone with modulus $\mu = 1/\lambda$. Stochastic approximation with step $1/(\mu j) = \lambda/j$ then attains $\mathbb{E} \| x^j - J_\lambda(z^m) \|^2 = O(\lambda^2 \sigma^2 / j)$, where σ^2 bounds the mini-batch variance (Koshal et al., 2013).

Fact 2: the step is capped by λ. Since $\gamma_j \le \lambda$ for every j, a large user-chosen γ_0 cannot push the iterates further than the proximal term allows. A small γ_0 matters only for the first λ/γ_0 inner steps of each round. The parameter that controls behaviour is λ, which has a direct meaning: the scale of movement allowed in one round. Choosing λ is easier than choosing γ_0, whose right value depends on the unknown size of F and on the variance factor $1/(1-\alpha)$.

Proposition 3 (outer convergence). Let Assumption M hold, let (4) have a solution, and let $e_m = z^{m+1} - J_\lambda(z^m)$ satisfy $\sum_m \mathbb{E} \| e_m \| < \infty$. Then z^m converges almost surely to a solution of (4).

Proof sketch. J_λ is firmly non-expansive because $F + N_X$ is maximal monotone. The exact iteration $z^{m+1} = J_\lambda(z^m)$ therefore converges to a zero of $F + N_X$, which is Rockafellar's theorem (1976, Thm. 1). Summable errors preserve convergence by the same theorem, criterion (A), applied pathwise. Integrability is ensured by the expectation condition and the monotone convergence theorem. ▫

Remark 3 (inner budget). By Fact 1, $\mathbb{E} \| e_m \| = O\left(\lambda\sigma/\sqrt{n_m}\right)$. The summability in Proposition 3 therefore needs n_m to grow faster than m^2, for example $n_m \propto m^{2+\delta}$. In our experiments we used the cheaper linear schedule $n_m = 20 + 4m$. It was enough in practice, but it is not covered by Proposition 3. A sample-complexity analysis of the linear schedule is left open.

Remark 4 (relation to other schemes). PGD is a stochastic, inexact instance of the proximal point method. It is close in spirit to the iterative Tikhonov regularisation of Koshal, Nedić and Shanbhag (2013), except that we re-centre the regulariser at each round instead of letting its weight decay. Extragradient-type methods such as stochastic mirror-prox (Juditsky et al., 2011) reach $O\left(1/\sqrt{K}\right)$ rates under Lipschitz continuity of F. Our F is discontinuous in t_i whenever L_i has atoms. Learning-in-games methods with bandit feedback (Mertikopoulos & Zhou, 2019) address a harder information setting than the full-gradient oracle assumed here.

5. Numerical study

5.1. Test market

The test market has $N = 4$ firms. The inverse demand is $p = a - bQ + \xi$ with $a = 20$, $b = 1$ and $\xi \sim \mathcal{N}(0, 3^2)$.

Firm i's loss is

$$L_i(q;\omega) = q_i\,(c_i - a + bQ - \xi) + \frac{1}{2}\,\gamma_i\,\zeta\,q_i^2,$$

with $c = (2,3,4,5)$ and $\gamma = (0.5,0.5,1,1)$. The cost shock is $\zeta = \exp(0.5z - 0.125)$, so that $\mathbb{E}[\zeta] = 1$, where z is standard normal with $\mathrm{corr}(z,\xi) = -0.6$. High input costs therefore tend to coincide with weak demand, which makes the left tail of profit heavy. Capacities are $\bar{q}_i = 20$ and $T = 300$.

The reference equilibria are computed on a fixed sample of 20,000 scenarios (seed 11). We use Gauss–Seidel best-response iteration, with each best response found by bounded scalar minimisation of the empirical CVaR. Accuracy is measured on the same sample by the Euclidean distance $\| q - q^* \|$ and by the Nikaido–Isoda gap $\sum_i [\rho_i(q) - \min_y \rho_i(y, q_{-i})]$.

5.2. What risk aversion does to the market

Table 1: Reference equilibria by CVaR level α. Expected profit is computed under the true distribution, that is, risk-neutrally. In every case the Nikaido–Isoda gap is below 10^{-12}.

α	q_1	q_2	q_3	q_4	Total Q	Mean price	Total expected profit
0 (risk-neutral)	4.237	3.569	2.178	1.677	11.660	8.340	49.65
0.5	3.639	3.027	1.740	1.299	9.705	10.295	62.24
0.9	2.895	2.364	1.248	0.882	7.389	12.611	66.31
0.95	2.658	2.154	1.103	0.760	6.675	13.325	65.19

Every firm produces less as α rises. The cut is proportionally largest for the high-cost firm 4, whose output falls by 47% between $\alpha = 0$ and $\alpha = 0.9$, against 32% for firm 1. The interesting effect is on profit. Moving from $\alpha = 0$ to 0.9 *raises* total expected profit by 34%. Each firm's fear of the tail restrains output, and that restraint moves the market toward lower, more profitable output without any coordination. Past $\alpha = 0.9$ the restraint begins to overshoot, and expected profit falls again (66.31 to 65.19).

5.3. Robustness to the step size

We fix $\alpha = 0.9$. Each method receives 20,000 oracle calls with mini-batches of $B = 16$, which is 320,000 scenarios in total, starting from $q = (1,1,1,1)$ and $t = 0$. PGD uses $\lambda = 0.5$ with the inner schedule $n_m = 20 + 4m$, which gives 96 outer rounds. SPG reports its last iterate. Each configuration is repeated 10 times with independent sample streams.

Table 2: Accuracy after 20,000 iterations at $\alpha = 0.9$: means over 10 runs, standard errors in parentheses.

Initial step γ_0	SPG $\| q - q^* \|$	SPG NI gap	PGD $\| q - q^* \|$	PGD NI gap
0.05	0.661 (0.002)	0.412 (0.005)	0.032 (0.002)	0.003 (0.001)
0.2	0.030 (0.005)	0.006 (0.003)	0.061 (0.007)	0.017 (0.005)
1.0	0.067 (0.018)	0.043 (0.016)	0.140 (0.010)	0.086 (0.014)
3.0	0.100 (0.023)	0.085 (0.035)	0.140 (0.010)	0.086 (0.014)
Worst over the four steps	0.661	0.412	0.140	0.086

The table supports a cautious reading.

• **PGD is more robust.** Over a 60-fold range of γ_0, its worst distance to equilibrium (0.140) is about a fifth of SPG's (0.661). The SPG failure occurs at the small step. With $\gamma_0 = 0.05$ and a $1/\sqrt{k}$ decay, its steps shrink before it arrives. It closes only about 70% of the initial distance of 2.35. PGD avoids this because each round is re-anchored at a point that is already close.

• **The cap is visible in the data.** The PGD results at $\gamma_0 = 1.0$ and 3.0 are identical. This is Fact 2: once γ_0 exceeds $\lambda = 0.5$, the schedule $\gamma_j = \min\{\gamma_0, \lambda/j\}$ no longer depends on γ_0.

• **Tuned SPG is more accurate.** At its best step ($\gamma_0 = 0.2$), SPG reaches 0.030 against PGD's 0.061 at the same step. This test game appears to be strongly monotone already, so the proximal term adds bias per round without adding needed curvature. PGD's best result, 0.032 at $\gamma_0 = 0.05$, matches SPG's best.
In short, PGD trades a little peak accuracy for not having to know the right step size. When α is a modelling choice that will be varied, and the variance factor $1/(1-\alpha)$ varies with it, that is usually the better trade.

6. Conclusion

CVaR CVaR reduces each agent's stochastic game problem to a minimax game against a bounded reweighing of nature. The Rockafellar-Uryasev variable t i transforms the collection of problems into a single variational inequality on an expanded space. Existence is straightforward by classical convex-game arguments. Computation is where the risk level plays a role: the stochastic pseudo-gradient has a variance of order $1/(1-\alpha)$, and a step size suitable for one level of risk-aversion is likely to be poor for a different level.
Proximal game dynamics include a quadratic "reluctance to move" to every agent each round. This strongly monotone, upper-bounded λ step makes each round invariant with λ. In our test market the cap reduced worst-case error across 60-fold steps from 0.661 to 0.140, while a carefully calibrated plain method remained twice as accurate. The economics side of the experiment is a reminder that risk aversion is public, at least in a Cournot market, where it reduces output and, to a limit, increases everyone's expected profit.
The main open issues are: (a) Proposition 3 required inner budgets growing faster than m2, whereas our experiment implemented a linear schedule, so a theory of cheaper schedule. (b)Monotonicity of the extended CVaR map was not established, and verifiable sufficient conditions would fully determine the self-contained framework. (c)Heterogeneous risk levels αi, and generalised (shared-constraint) equilibria (Facchinei & Kanzow, 2007), are natural next steps.
Data and code availability. All data are simulated. The Python code for reproducing Tables 1 and 2 (seeds 11, 200-209) can be obtained from the author.
Competing interests. The author has declared no conflict of interest.

References

1. Artzner, P., Delbaen, F., Eber, J. M., & Heath, D. (1999). Coherent measures of risk. *Mathematical finance*, *9*(3), 203-228. https://doi.org/10.1111/1467-9965.00068
2. Facchinei, F., & Kanzow, C. (2010). Generalized Nash equilibrium problems. *Annals of Operations Research*, *175*(1), 177-211. https://doi.org/10.1007/s10288-007-0054-4
3. Juditsky, A., Nemirovski, A., & Tauvel, C. (2011). Solving variational inequalities with stochastic mirror-prox algorithm. *Stochastic Systems*, *1*(1), 17-58. https://doi.org/10.1287/10-SSY011
4. Koshal, J., Nedic, A., & Shanbhag, U. V. (2012). Regularized iterative stochastic approximation methods for stochastic variational inequality problems. *IEEE Transactions on Automatic Control*, *58*(3), 594-609. https://doi.org/10.1109/TAC.2012.2215413
5. Mertikopoulos, P., & Zhou, Z. (2019). Learning in games with continuous action sets and unknown payoff functions. *Mathematical Programming*, *173*(1), 465-507. https://doi.org/10.1007/s10107-018-1254-8
6. Nash, J. F. (1950). Equilibrium points in n-person games. *Proceedings of the national academy of sciences*, *36*(1), 48-49. https://doi.org/10.1073/pnas.36.1.48
7. Ravat, U., & Shanbhag, U. V. (2011). On the characterization of solution sets of smooth and nonsmooth convex stochastic Nash games. *SIAM Journal on Optimization*, *21*(3), 1168-1199. https://doi.org/10.1137/100792644
8. Rockafellar, R. T. (1976). Monotone operators and the proximal point algorithm. *SIAM journal on control and optimization*, *14*(5), 877-898. https://doi.org/10.1137/0314056
9. Rockafellar, R. T., & Uryasev, S. (2000). Optimization of conditional value-at-risk. *Journal of risk*, *2*, 21-42. https://doi.org/10.21314/JOR.2000.038
10. Rosen, J. B. (1965). Existence and uniqueness of equilibrium points for concave n-person games. *Econometrica: Journal of the Econometric Society*, 520-534. https://doi.org/10.2307/1911749
11. Azizov Azizbek Sabitkhanovich. (2026). Phase Space Analysis of Multi-Scale Nonlinear Epidemic Systems with Behavioral Feedback. *Applied Nonlinearity in Science and Technology*, *2*(3), 16–20.
12. Kumar, M. S. (2026). Mathematical Formulation with Finite Element Simulation for Thermo-Mechanical Stress Analysis in Power Electronic Modules. *Journal of Applied Mathematical Models in Engineering*, 22-33.
13. Lemeon, M., Regash, J., & Leyene, T. (2026). Mathematical Formulation for Dynamic Simulation of Grid-Forming Virtual Synchronous Generator Control under Low-Inertia Conditions. *Journal of Applied Mathematical Models in Engineering*, 42-51.
14. Rajan, C. (2026). Congestion-and fault-severity-aware escape routing with local recovery for multicore networks-on-chip. *VSF Journal of Electronics and Communication Engineering*, *1*(1), 30-37.
15. Rane Kuma, & Sikalu T C. (2026). Quantum-Resilient Cryptographic Implementations for Embedded Communication Systems. *Progress in Electronics and Communication Engineering*, *3*(1), 1-11.
16. Navamani, G., Sumathi, N., & Karthikeyan, K. R. (2026). Exploring the certified bondage number in

diverse types of graphs. *Results in Nonlinear Analysis*, *9*(3), 106-114.
17. Eliana Marbel Rivas Téllez. (2026). Coupled Thermo-Mechanical Finite Element Analysis of Additively Manufactured Lattice Structures under Cyclic Loading Conditions. *Advances in Mechanical Engineering and Applications*, *2*(3), 25-32.
18. J.Karthika. (2026). Multi-Objective Optimization of Hybrid Hydrogen–Battery Energy Management for Heavy-Duty Electric Vehicles Using NSGA-II. *Advances in Mechanical Engineering and Applications*, *2*(3), 1-11.
19. P Kalaivanai. (2026). Assistive Intelligent Communication Models for Peer-Based Online Learning Environments. *Journal of Intelligent Assistive Communication Technologies*, *2*(1), 72-80.