Resource Allocation in Multi-Domain Operations: A Multi-Resource Colonel Blotto Game Approach

Rogério Barbosa Marques[1], Raphael de Almeida Leitão[1], Fábio Ayres Cardoso[1]
[1,]Universidade da Força Aérea (UNIFA), Rio de Janeiro/RJ – Brasil

***Abstract* – Multi-domain operations are an emerging doctrinal concept that proposes the simultaneous and interdependent employment of forces across land, sea, air, space, and cyberspace to create windows of advantage against peer competitors using anti-access/area denial strategies. This study asks whether spreading scarce military resources across additional domains is advantageous relative to the classic joint allocation over land, sea, and air. It was modeled into two antagonistic joint forces as a two-person zero-sum Multi-Resource Colonel Blotto game, in which each operational domain is a battlefield. Mixed-strategy Nash equilibria are computed using a fictitious-play procedure. The results show that, against an adversary confined to three classic domains, a force distributing resources across five domains attains an equilibrium win probability that rises with the resource budget and settles near 0.68, whereas symmetric doctrines tie. These findings support the doctrinal value of fractionating forces across domains.**

***Keywords* – Colonel Blotto game, multi-domain operations, game theory.**

I. INTRODUCTION

Every five years, the North Atlantic Treaty Organization (NATO) re-examines its strategic concept. The 2022 concept stresses the holistic impact of multi-domain operations across a variety of conflicts, not only between great powers [1]. A domain is "any potential operational space through which the target system can be influenced — not only the land, sea, air and space domains, but also the virtual (information and cyber) and the human (cognitive, moral and social)" [2, p. 16]. It can also be understood as the space "where forces employ their capabilities and conduct operations" [3].

Future wars "will require more than a joint mindset or a synchronized joint force; they must also consider a force that conducts interdependent operations" [4], whose joint and simultaneous employment across several domains creates windows of advantage to seize the initiative from potential adversaries and restore credible conventional deterrence against peer competitors. Multi-domain operations are an emerging doctrinal concept developed in the United States that proposes solutions to confront adversaries with layered, long-range (stand-off) weapon capabilities, principally China and Russia, which have built anti-access/area-denial (A2/AD) strategies that separate the joint force in time, space, and function.

Given the importance of these operations in future conflicts and the scarcity of resources to be used across all domains, one motivating question is whether allocating resources to additional domains (space, cyberspace, and others), beyond the classic joint domains (land, air, and sea), offers an advantage over the current doctrine.

Rogério Barbosa Marques, marquesppgao@gmail.com; Raphael de Almeida Leitão, leitao_raphael@hotmail.com; Fábio Ayres Cardoso cardosoaf@yahoo.com.br.

Military exercises and war games are one way to study this question, but they are costly and difficult to run iteratively. This study uses game theory, specifically the Colonel Blotto framework, to study whether it is advantageous to allocate resources across the operational domains that compose multi-domain operations. Although a mathematical model cannot capture every variable of a real problem, it is useful for approximating an interaction and understanding the underlying processes [5]. One delimitation should be stated at the outset, before the operational reader turns to the results: the model developed here treats the operational domains as linearly independent, so cascading cross-domain failures are not reflected in the win probabilities reported below.

The qualitative case for this shift is well-established in doctrinal literature. Marques et al. [17], reviewing the impact of multi-domain thinking on aerospace power, characterized it as a broad way of thinking about the convergent employment of military and non-military means across the classic domains augmented by space, cyberspace, the electromagnetic, and the cognitive (human) domains, so as to overcome vulnerabilities that would arise if capabilities were applied separately and to create dilemmas that degrade an opponent's will to react. They concluded that current and future conflicts tend to occur in multi-domain environments, which reshape command-and-control structures, joint employment doctrine, and the pursuit of aerospace superiority. The present work complements that qualitative analysis by asking, in quantitative terms, whether and by how much distributing resources across additional domains pays off.

II. THE MULTI-RESOURCE COLONEL BLOTTO GAME

The Colonel Blotto game is a classic game-theoretic problem named after a deadlock faced by the legendary Colonel Blotto, who had to divide his attacking force among many forts without knowing how the defending forces were distributed among them. Each player divides a force across many areas; the payoff is the sum of the area payoffs, and each area payoff depends only on the force allocated there [6], [7]. Because the game is closely related to military resource-allocation problems, it provides an ideal framework for analyzing whether multi-domain operations are more suitable than classic joint operations.

Because the proposed problem involves players (antagonistic forces) with resources that can be distributed across several fronts (domains) and battlefields (theaters of operation), it is analyzed using what has been termed a Multi-Resource Blotto game [9]. Players assign platforms (ships, aircraft, tanks, etc.) over a set of battlefields. Each platform carries a set of capabilities, and each unit of capability carried by a platform is a resource. The winner of a battlefield is a

function of the joint allocation of resources, and the goal is to win as many battlefields as possible [14].

As an illustration, consider two commanders who must split their resources between two theaters. Each must allocate 10 ships and 15 aircraft to the task force. Each ship carries five missiles and a radar, and each aircraft carries two missiles, a radar, and a camera. Ships and aircraft are platforms, and missiles, radars, and cameras are the capabilities. The theater winner is a function of the joint allocation of missiles, radars, and cameras; for example, "the player who assigns more missiles, provided at least two radars and one camera are assigned, wins" [14]. Fig. 1, for example, illustrates this situation, where we can see that a commander divides naval and air platforms — each carrying missile, radar, and camera capabilities — between two theaters of operations.

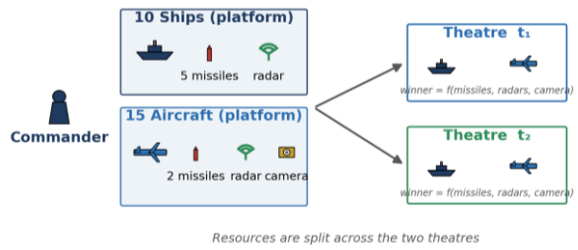

Fig. 1. Illustration of the Multi-Resource Blotto game: a commander splits ship and aircraft platforms (each carrying missiles, radar, and camera capabilities) across two theaters of operation.

III. RELATED WORK

The Colonel Blotto game originates from the continuous formulation of Gross and Wagner [6], who derived equilibrium allocations for symmetric and asymmetric budgets. Roberson [7] characterized the unique equilibrium payoffs of the continuous game with homogeneous battlefields and showed that, at equilibrium, each player's allocation to any battlefield follows a uniform marginal distribution, which is the formal counterpart of the hedging behavior discussed in Section VI. Hart [8] developed discrete and General Lotto variants, which are closer to the integer-resource setting relevant to military platforms.

On the computational side, the game is notoriously difficult because the strategy space grows combinatorially. Behnezhad et al. [9] gave a polynomial-time algorithm for optimal strategies of a broad class of Blotto games, and Ahmadinejad et al. [10] reduced equilibrium computation to a polynomial-size linear program via a clever bijection between battlefield allocations and bipartite matchings. However, for the multi-resource extension used here, exact methods remain expensive, which motivates the iterative fictitious-play approach of Lee [11], building on the convergence results of Brown [12] and Robinson [13].

The doctrinal relevance of dividing forces is supported by recent operations research studies. Lamb, Grana, and O'Donoughue [14] analytically and through simulation show that fractionating forces across more battlefields improves competitive outcomes, and Kovenock and Roberson [18] establish that asymmetries in the number of contested dimensions translate directly into payoff asymmetries, a result later consolidated in their survey of conflicts with multiple battlefields [15]. Gulden and Grana [16] reach compatible conclusions on mosaic warfare using the Blotto model. The results of this study reproduce and quantify these effects in the specific context of the five operational domains of multi-domain operations.

IV. MODEL

There are two players, A and B. Let $C = \{c_1, c_2, \ldots, c_M\}$ be the set of capabilities and $P_i = \{p_{i1}, p_{i2}, \ldots, p_{iKi}\}$ the platforms of player i. Each platform is an integer vector whose k-th component is the number of capability-k resources it carries. The number of battlefields is T. Players simultaneously distribute all their resources over the battlefields. Fig. 2 sketches the structure of the model, mapping each player's platforms and capabilities onto the five operational domains treated as battlefields. In this case, players A and B allocate the resources carried by their platforms (via the allocation $x_{j,t}$) across the battlefields/domains; each domain yields a payoff $U_{j,t}$ summed into W_i.

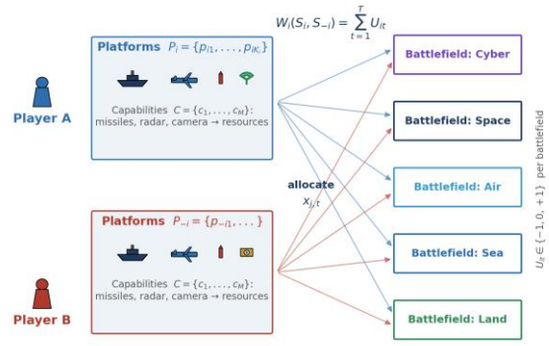

Fig. 2. Structure of the Multi-Resource Blotto model: players A and B allocate the resources carried by their platforms (via the allocation $x_{j,t}$) across the battlefields/domains; each domain yields a payoff $U_{j,t}$ summed into W_i.

A pure strategy partitions a player's resources across battlefields. Let x and y be the allocation vectors of A and B. For a given strategy, let $Q_t^{S_i}$ denote the number of resources of each capability player i assigns to battlefield t. Given the joint allocations, the payoff of player i is the sum of the per-battlefield payoffs:

$$W_i(S_i, S_{-i}) = \sum_{t=1}^{T} U_{it}\left(Q_t^{S_i}, Q_t^{S_{-i}}\right) \tag{1}$$

In the simple Blotto rule, U_{it} equals +1 if player i assigns more resources than $-i$ to battlefield t, 0 if equal, and −1 otherwise. Because the game is zero-sum, $U_{-it} = -U_{it}$. A Nash equilibrium is a strategy profile in which both players maximize the expected utility W, with the expectation being taken over the mixed-strategy distributions.

To make the payoff concrete, consider three battlefields $T = \{t_1, t_2, t_3\}$ where each per-battlefield outcome $U_{it} \in \{-1, 0, +1\}$ is equally probable. The total payoff for player A, $W_A = U_{A,1} + U_{A,2} + U_{A,3}$, then ranges over $\{-3, \ldots, +3\}$ with the multinomial distribution of Table I, over $n = 27$ equally likely combinations. The expected value $E(W_A)$ is zero, confirming

that, in the absence of doctrinal asymmetry, neither player is favored (the baseline against which the advantages reported in Section VI are measured).

TABLE I. DISTRIBUTION OF THE TOTAL PAYOFF W_A

W_A	−3	−2	−1	0	+1	+2	+3
$P(W_A = w)$	1/27	3/27	6/27	7/27	6/27	3/27	1/27

From the equilibrium, the paper derives the quantity reported in Figs. 4 and 5, the win probability ρ_A of player A. Let $W^* = (x^*)^T M y^*$ be the value of the game at the equilibrium mixed strategies (x^*, y^*), where M is the payoff matrix and $-T \leq W^* \leq T$. Because each of the T battlefields contributes a payoff in {−1, 0, +1}, the normalized value W^*/T lies in [−1, 1]; mapping it onto [0, 1] yields the win probability:

$$\rho_A = \frac{1}{2}(W^*/T) + \frac{1}{2} \tag{2}$$

Thus, $\rho_A = 1/2$ corresponds to a balanced contest $W^* = 0$, $\rho_A = 1$ to certain victory ($W^* = T$), and $\rho_A = 0$ to certain defeat. In Fig. 4, the equilibrium mixed strategy is summarised as the share of the budget allocated to each domain by player A. In Fig. 5, the scalar ρ_A of (2) is plotted against the total resources for each doctrine pairing.

The numerical experiments assume symmetric resource budgets for both players. It is important to highlight that this is a controlled simplification, since real A2/AD scenarios tend to exhibit asymmetries.

V. SOLUTION METHOD

Even modest instances generate enormous strategic spaces. The number of pure strategies of a player allocating "f" resources over n battlefields equals the number of compositions $c(n + f - 1, f)$, which grows combinatorially. Fig. 3 shows this growth on a logarithmic scale. For $f = 20$ resources, the count already reaches 2.8×10^{13} strategies at $n = 30$ domains, which rules out the direct enumeration of the payoff matrix for realistic sizes.

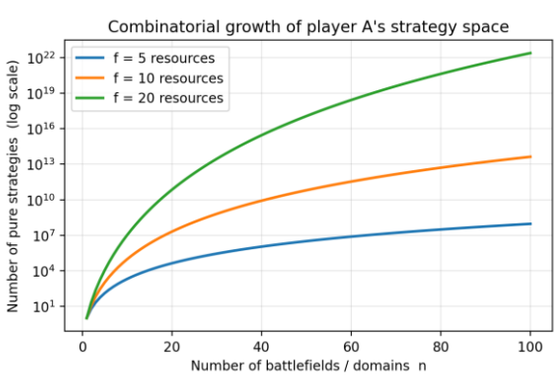

Fig. 3. Combinatorial growth of the number of pure strategies, C(n + f − 1, f), with the number of domains n and resources f.

To obtain mixed-strategy equilibria, this study uses a fictitious-play (FP) procedure [11], [12], [13]. FP is an iterative method in which, at each iteration, each player chooses the pure strategy that is the best response to the empirical mix of the opponent's past plays. It is known to converge to the optimal solution of two-person zero-sum games and is often the only realistic approach for large games [11]. The value of the game and the equilibrium mixed strategies are recovered from empirical frequencies of play.

In the present implementation the payoff matrix of the instance is built explicitly and the Brown–Robinson iteration is run for a fixed number of steps, initialised at the maximin row and the minimax column and without a convergence test; the value of the game is then read from the empirical frequencies of play. Two deliberate simplifications keep the computation tractable. First, the instance solved is the single-resource, single-platform specialization of the model of Section IV, in which the capability vector of a platform collapses to a scalar count of fungible resources and the platform layer of Section IV is not instantiated, so that a pure strategy is a composition of the total budget over the domains. Second, whenever the number of compositions exceeds the cap reported in Table II, the strategy set of each player is replaced by a uniform random sample of that size drawn with a fixed seed; the reported equilibria are therefore equilibria of a sampled subgame. The complete parameterization is given in Table II so that the experiments can be reproduced, and the scope of both simplifications is revisited in Section VII.

VI. RESULTS

A. Equilibrium allocation

In the symmetric instance, where both players contest the same five domains with the same budget, the fictitious-play procedure converges to a game value close to zero, as expected for a fair contest, and the equilibrium marginal allocation is uniform across the domains to within numerical precision: each domain receives 0.20 of the budget, with a residual spread below 0.01 that shrinks as the iteration count grows. The equilibrium therefore spreads probability mass broadly rather than concentrating on a single domain: a hedging behavior that is the game-theoretic signature of the Blotto structure.

This diffuse allocation is the discrete-game counterpart of Roberson's result [7], which states that, in the continuous symmetric game, each battlefield receives a uniform marginal allocation. Intuitively, any deterministic concentration of force would be exploited by an opponent who shifts mass to the abandoned domains; randomization removes this vulnerability. The practical reading is that no single domain should be treated as decisive *a priori* — the optimal posture keeps the adversary uncertain about where strength will fall, which is precisely the "windows of advantage" logic of the multi-domain doctrine. The asymmetric contest is more informative, and Fig. 4 reports it. Against an adversary confined to the three classic domains, the multi-domain force keeps the bulk of its budget on the three shared domains and sends only a small residue to space and cyberspace, evenly split between the two, and that residue shrinks as the budget grows. The reason is that prevailing against an allocation of zero takes a single unit whatever the size of the force: holding both uncontested domains costs at least 2/f of the budget. This

is the mechanism behind the win probabilities of Fig. 5: two domains are secured at a cost that becomes negligible as the force grows, and the three contested domains are then disputed on nearly equal terms. Note that the quantity plotted here is a share of the budget, not the scalar win probability of (2).

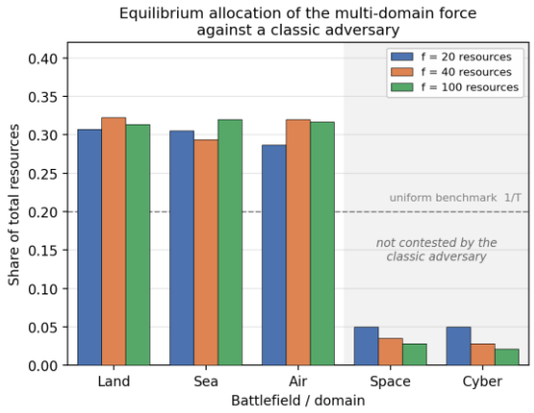

Fig. 4. Equilibrium marginal allocation of the multi-domain force (player A) against a classic adversary, by domain and total resource budget. The dashed line marks the uniform share $1/T$ of the symmetric contest.

B. Effect of doctrine and resources

The central question is whether distributing resources across the five multi-domain battlefields (land, sea, air, space, and cyberspace) outperforms the classic three-domain allocation (land, sea, air). Fig. 5 reports the equilibrium win probability of player A as a function of the total resources for the four doctrine pairings. When a multi-domain force faces an adversary confined to the three classic domains, its win probability rises from 0.60 at the smallest budget to a plateau of about 0.68, which is reached once each player holds roughly forty resources. The reverse pairing (classic against multi-domain) falls symmetrically to about 0.32, whereas the two symmetric pairings coincide at 0.5, as the zero-sum structure requires.

For reproducibility, Table II summarizes the numerical instance used to generate the results in Fig. 5. Both players were assigned symmetric resource budgets, while the number of contested domains varied according to the doctrine pairing.

TABLE II. NUMERICAL EXPERIMENT CONFIGURATION

Parameter	Value / Configuration
Players	2 (A and B)
Classic doctrine	3 domains: Land, Sea, Air
Multi-domain doctrine	5 domains: Land, Sea, Air, Space, Cyber
Platforms per player (K)	1 (platform layer not instantiated)
Number of capabilities (M)	1 (single fungible resource type)
Budgets shown in Fig. 4	20, 40 and 100 resources
Total resources per player	{6, 10, 20, 40, 70, 100}
Resource budgets	Symmetric between players
Payoff rule	Simple majority: +1, 0, or -1 per contested domain
Doctrine pairings	Classic–Classic; Classic–Multidomain; Multidomain–Classic; Multidomain–Multidomain
Solution method	Fictitious-play
Pure strategies per player	Full set if ≤ 5,000; otherwise a uniform random sample of 5,000 (fixed seed)
FP iterations	1,200
Convergence / stopping criterion	None; fixed iteration count

For each resource level, the fictitious-play procedure was applied to the four doctrine pairings shown in Table II. The resulting equilibrium game value was converted into the win probability ρ_A according to (2).

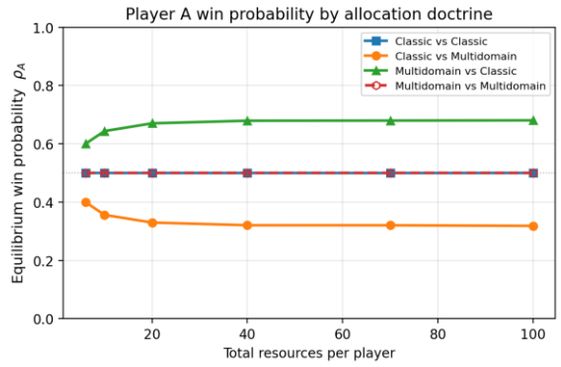

Fig. 5. Equilibrium win probability ρ_A of player A by allocation doctrine, as a function of total resources per player.

The advantage is structural rather than a matter of force size: it is present at every budget tested and does not require numerical superiority, indicating that the benefit comes from contesting additional, otherwise uncontested, domains. This is consistent with the literature on fractionation, which finds that dividing forces across more battlefields improves competitive outcomes when the opponent cannot match the dimensionality of the contest [14], [15], [16].

Three observations are made. First, the gap is already evident at small resource totals and does not require numerical superiority, which suggests that the dimensionality of the contest is a lever distinct from mass. Second, the advantage widens as the budget grows and then flattens: a larger budget lets the multi-domain force hold the two uncontested domains while still competing for the three shared ones, and once that division of effort is affordable, further resources add little. The doctrinal choice of how many domains to contest therefore dominates the marginal-resource choice from about forty resources onwards. Third, the symmetry between the classic-versus-multidomain and multidomain-versus-classic curves around 0.5 reflects the zero-sum structure: the advantage gained by one doctrine is exactly the disadvantage imposed on the opponent, so unilaterally adopting the broader doctrine is a dominant move whenever the adversary has not yet done so.

VII. DISCUSSION

Five modeling choices shape the interpretation of these results. First, the simple Blotto payoff rewards majority allocation per domain; richer rules in which winning a domain requires prerequisite capabilities — the radar-and-camera condition of the illustrative example — would penalize forces that fractionate so finely that they fail to meet capability thresholds, tempering the advantage of the broadest doctrine. Second, the model treats the five domains as homogeneous and independent, whereas in practice, space and cyberspace effects are enabling and cross-cutting, so a win in one domain can change the payoff structure of another; capturing these couplings would require domain-specific and possibly non-separable payoff functions [15].

Third, the contest is symmetric in force size. Asymmetric budgets, which are the realistic case against A2/AD-equipped peers, are known to shift equilibrium payoffs non-trivially [6], [7], and an under-resourced force that nonetheless contests more domains may or may not retain the advantage observed here.

Fourth, the combinatorial formulation assumes that operational resources are perfectly fungible across domains, that is, it abstracts the vectors as generic, domain-agnostic efforts. In tactical reality, however, physical platforms are subject to rigid allocation constraints, since naval, land, air, space, and cyber assets cannot be freely reassigned to one another. Imposing domain-compatibility constraints would drastically reduce the space of pure strategies and would yield a more realistic representation of force allocation.

Fifth, the figures rest on two computational approximations. The instance solved is the single-resource, single-platform specialisation of the model, so neither the prerequisite-capability rules discussed above nor any trade-off in assigning a given platform to a given domain can bind, and for the larger budgets the strategy sets are uniform random samples of the composition space rather than the space itself, which makes the reported values equilibria of a sampled subgame. The sample size was calibrated against exact enumeration, which is feasible up to forty resources per player: over that range the sampled and the exact values agree to within 0.005, and the residual bias is negative, so the advantage reported for the multi-domain doctrine is a conservative estimate.

The same restriction explains why the shares that Fig. 4 assigns to the uncontested domains sit above the 2/f floor at the larger budgets: the sampled strategy set still contains allocations more generous to those domains than the equilibrium requires. The polynomial-time formulations of [9] and [10] would remove this restriction altogether. Instantiating the platform layer, so that a pure strategy becomes a joint assignment of K platforms over T domains, is the natural next step and would allow the compatibility constraints raised above to be imposed directly.

These caveats do not overturn the central finding — contesting additional domains is advantageous against an opponent that cannot — but they delimit its scope and identify the parameters that most deserve sensitivity analysis.

Operationally, the results argue for treating domain breadth as a first-order planning variable, on par with force mass. They also reinforce the value of fictitious-play as a tractable tool: because it scales to games far beyond the reach of explicit payoff-matrix methods [11], it can support iterative wargame-style exploration of doctrinal options at a fraction of the cost of live exercises.

VIII. CONCLUSION

Modeling two joint forces as a Multi-Resource Colonel Blotto game and solving it by fictitious-play, this paper finds quantitative support for the doctrinal claim behind multi-domain operations: allocating resources across the additional space and cyberspace domains yields a clear equilibrium advantage over an adversary that remains in the three classic domains, and this advantage is robust to the total volume of resources. The result echoes findings on fractionation and mosaic warfare [14], [16], suggesting that the dimensionality of the contest, not only the mass of forces, is a decisive variable.

It is important to highlight that the value 0.68 is obtained using the simple majority rule of the model and may be optimistic, considering that a highly fragmented force may not achieve the minimum capabilities necessary to operate in a given domain (minimum force packages).

Future work should incorporate asymmetric forces, heterogeneous capabilities, and domain-specific payoff functions to refine these estimates.

REFERENCES

[1] NATO, “NATO 2022 Strategic Concept”, North Atlantic Treaty Organization, Madrid, June 2022.
[2] U.S. Department of Defense, “Capstone Concept for Joint Operations — Version 2.0”, Washington, DC, August 2005.
[3] R. B. Marques, M. D. O. Dias, F. N. Camargo, “Literature review on multi-domain battle”, Int. J. of Development Research, vol. 12, no. 4, pp. 55380–55384, 2022.
[4] D. A. Deptula, “Interdependent warfare: combined effects power in the 21st century”, Mitchell Institute for Aerospace Studies, vol. 10, March 2018.
[5] G. E. P. Box, “Robustness in the strategy of scientific model building”, in Robustness in Statistics, Academic Press, pp. 201–236, 1979.
[6] O. Gross, R. Wagner, “A continuous Colonel Blotto game”, RAND Corporation, RM-408, Santa Monica, 1950.
[7] B. Roberson, “The Colonel Blotto game”, Economic Theory, vol. 29, no. 1, pp. 1–24, 2006.
[8] S. Hart, “Discrete Colonel Blotto and General Lotto games”, Int. J. of Game Theory, vol. 36, pp. 441–460, 2007.
[9] S. Behnezhad, S. Dehghani, M. Derakhshan, M. T. HajiAghayi, S. Seddighin, “Faster and simpler algorithm for optimal strategies of Blotto game”, 31st AAAI Conf. on Artificial Intelligence, pp. 369–375, 2017.
[10] A. Ahmadinejad, S. Dehghani, M. Hajiaghayi, B. Lucier, H. Mahini, S. Seddighin, “From duels to battlefields: computing equilibria of Blotto and other games”, Mathematics of Operations Research, vol. 44, no. 4, pp. 1304–1325, 2019.
[11] M. G. Lee, “New fictitious play procedure for solving Blotto games”, M.S. thesis, Naval Postgraduate School, Monterey, CA, December 2004.
[12] G. W. Brown, “Iterative solution of games by fictitious play”, in Activity Analysis of Production and Allocation, Wiley, pp. 374–376, 1951.
[13] J. Robinson, “An iterative method of solving a game”, Annals of Mathematics, vol. 54, no. 2, pp. 296–301, 1951.
[14] J. Lamb, J. Grana, N. O’Donoughue, “The benefits of fractionation in competitive resource allocation”, Computational Economics, vol. 59, no. 2, pp. 831–852, 2022.
[15] D. Kovenock, B. Roberson, “Conflicts with multiple battlefields”, in Oxford Handbook of the Economics of Peace and Conflict, Oxford Univ. Press, 2012.
[16] T. Gulden, J. Grana, “Findings on mosaic warfare from a Colonel Blotto game”, RAND Corporation, RR-A1611, Santa Monica, 2021.
[17] R. B. Marques, L. F. S. Lima, G. W. O. Silva, R. A. Leitão, “Multi-domain thinking and its impacts on Aerospace Power”, Revista da UNIFA, Rio de Janeiro, vol. 38, pp. 1–36, 2025.
[18] D. Kovenock, B. Roberson, “Asymmetric conflicts with endogenous dimensionality”, Korean Economic Review, vol. 26, no. 2, pp. 287–305, 2010.