TECHNICAL UNIVERSITY OF CLUJ-NAPOCA

ACTA TECHNICA NAPOCENSIS

Series: Applied Mathematics, Mechanics, and Engineering
Vol. 69, Issue Special III, July, 2026

DYNAMIC BAYESIAN DECISION MODELING FOR AUTONOMOUS UAV ENGAGEMENT UNDER CONDITIONS OF UNCERTAINTY

Emanuelle PERȚA, Alina TOTOI

Abstract: *Autonomous unmanned aerial vehicles (UAVs) operate under significant uncertainty, where engagement decisions involve asymmetric and time-dependent costs. This paper proposes a Bayesian decision-theoretic framework for autonomous UAV engagement, explicitly modeling uncertainty, delay penalties, and abort options within an expected utility formulation. The decision problem considers three actions: engage, wait, and abort. It derives optimal policies as functions of target legitimacy probability and cost parameters. A numerical analysis based on synthetic simulations explores the decision space and identifies distinct operational regimes. The results demonstrate that delaying action can be optimal over broad regions of the parameter space, particularly under high uncertainty or moderate delay costs, highlighting the importance of risk-aware and temporally adaptive decision-making in autonomous UAV systems.*

Keywords: *Bayesian decision theory, Autonomous UAVs, Decision-making under uncertainty, Cost-sensitive decision models, Risk-aware systems.*

1. INTRODUCTION

Autonomous systems increasingly operate in environments characterized by uncertainty, incomplete information, and high-stakes consequences. In such contexts, decision-making cannot rely solely on deterministic rules or heuristic thresholds, as erroneous actions may incur severe operational, political, or societal costs. This challenge is particularly evident in autonomous unmanned aerial vehicle (UAV) operations, where engagement decisions must balance uncertain target identification against asymmetric costs associated with false engagement, delayed action, or mission abort. Designing decision mechanisms that are both mathematically sound and operationally interpretable remains an open problem in autonomous systems engineering.

From a theoretical perspective, decision-making under uncertainty has long been studied within the framework of Bayesian statistical decision theory. Classical works such as [1] and [2] establish expected utility maximization as the rational criterion for optimal decisions when system states are probabilistic and costs are asymmetric. These foundations have been widely applied across statistics, economics, and control theory, providing a rigorous basis for modeling risk-sensitive decisions. Within this framework, a decision-maker selects the action that maximizes expected utility, conditional on the available information and the associated loss or benefit structure.

Beyond static decision problems, many real-world systems require decisions to be made sequentially, where postponing action may allow the acquisition of additional information at the expense of time or operational cost. This perspective is formalized in sequential decision-making and optimal stopping theory, originally developed in the seminal work of Wald [3] and later extended through dynamic programming approaches [4]. These frameworks demonstrate that delaying a decision can be optimal when the expected value of additional information outweighs the cost of waiting. However, despite their mathematical maturity, such principles are

rarely made explicit in the design of autonomous engagement decisions, where systems are often constrained to immediate binary actions.

In the context of autonomous UAVs, existing research has primarily focused on coordination architectures, role assignment, and learning-based control policies. For instance, [5] proposes a Bayesian framework for multi-agent coordination under uncertainty, emphasizing situational awareness and role differentiation rather than explicit engagement thresholds. Similarly, survey studies such as [6] review decision-making mechanisms integrated into trajectory planning and information gathering. While these approaches demonstrate the relevance of probabilistic reasoning in autonomy, they often rely on complex architectures or learned policies that obscure the underlying decision logic and provide limited analytical insight into when action, delay, or disengagement is optimal.

Recent studies have also explored probabilistic decision-making in multi-UAV systems operating under uncertainty. For example, [7] propose a metacognitive decision-making framework for target search without inter-agent communication, emphasizing adaptive reasoning under uncertainty. While such approaches demonstrate the relevance of probabilistic decision models in UAV applications, they do not explicitly characterize engagement thresholds or delayed decision regions, which remain the focus of the present study.

As a result, a gap remains between classical decision-theoretic principles and their transparent application to autonomous engagement scenarios with asymmetric costs and uncertainty. In particular, few studies explicitly model the option to delay action as a first-class decision alternative alongside engagement and abort, nor do they provide a systematic characterization of decision regions as functions of probability and cost parameters. This lack of interpretability complicates both system verification and human oversight, especially in safety-critical environments.

The present paper addresses this gap by proposing a dynamic Bayesian decision-theoretic model for autonomous UAV engagement under uncertainty. The problem is formulated using expected utility theory, considering three possible actions: immediate engagement, delayed decision, and mission abort. The model explicitly incorporates asymmetric costs, a delay penalty, and probabilistic uncertainty regarding target legitimacy. The contribution lies in structuring and analyzing these classical principles within a unified and interpretable framework tailored to autonomous systems.

To investigate the behavior of the proposed model, a systematic numerical analysis is conducted using synthetic simulations. The parameter space is discretized over target legitimacy probability, relative cost ratios, and delay penalties, allowing the identification of distinct operational regimes. The resulting decision maps provide a transparent visualization of when engagement, delay, or abort becomes optimal, highlighting the dominant role of non-action and delayed decision-making in high-risk environments. By bridging established decision theory with an explicit and interpretable application to autonomous UAV engagement, this work contributes a rigorous yet practical perspective on risk-aware autonomy.

2. MATHEMATICAL MODEL AND METHODOLOGY

2.1 Decision-theoretic framework

The proposed approach is grounded in classical Bayesian decision theory, where decisions are made under uncertainty by maximizing expected utility. Let θ denote the true but unobserved state of the environment, representing the legitimacy of a detected target, with $\theta \in \{legitimate, non-legitimate\}$.

Based on sensor observations and prior information, the autonomous system maintains a belief:

$$p = P(\theta = legitimate \mid \mathcal{I}) \quad (1)$$

where $\mathcal{I}$ denotes the information available at the decision time. This probabilistic representation

of uncertainty follows standard formulations in Bayesian statistical decision theory [1,2].

The decision problem is defined over a discrete action space:

$$\mathcal{A} = \{STRIKE, WAIT, ABORT\} \quad (2)$$

corresponding to immediate engagement, delayed decision-making, and mission termination, respectively. The optimal action is selected according to the maximum expected utility criterion, consistent with the axiomatic foundations of rational decision-making [8].

2.2 Utility structure and cost asymmetry

To capture the asymmetric consequences of engagement decisions, a utility function $U(a, \theta)$ is defined for each action $a \in \mathcal{A}$. Let $B > 0$ denote the benefit associated with correctly engaging a legitimate target, and let $C > 0$ represent the cost incurred when engaging a non-legitimate target. These quantities may encompass operational losses, collateral damage, or broader strategic consequences.

For the STRIKE action, the expected utility is given by:

$$EU_{STRIKE}(p) = p \cdot B - (1 - p) \cdot C \quad (3)$$

which is the standard form for binary-state decisions with asymmetric costs [9].

The ABORT action yields a baseline utility $EU_{ABORT} = 0$, serving as a reference corresponding to risk avoidance through non-engagement.

In the absence of a delay option, immediate engagement would be optimal whenever

$$EU_{STRIKE}(p) > EU_{ABORT}, \quad (4)$$

leading to the classical decision threshold:

$$p > \tau = \frac{C}{B + C} \quad (5)$$

This threshold separates engagement and non-engagement regions and is a well-known result in cost-sensitive decision analysis [10].

In the absence of a delay option (WAIT), the decision problem reduces to a static comparison between engagement and abort. In this case, immediate engagement becomes optimal only when the probability of target legitimacy exceeds the cost-dependent threshold from (5). Figure 1 illustrates the behavior of this theoretical strike threshold as a function of the relative cost ratio $r = C/B$. As the engagement cost increases relative to the expected benefit, the required confidence level for engagement approaches unity, indicating an increasingly conservative decision policy.

This analytical threshold serves as a benchmark for the numerical results presented in the following section, where the introduction of a delay option fundamentally alters the decision structure

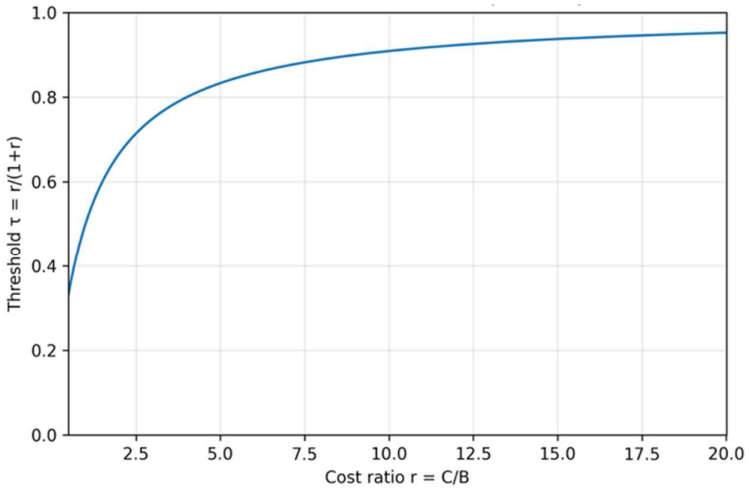


Fig. 1. Theoretical strike threshold as a function of the cost ratio $r = C/B$, without a delay option.

2.3 Modeling delayed decisions

Realistic autonomous systems often face situations where immediate commitment is suboptimal and additional information can be acquired at the expense of time or resources. To model this possibility, the WAIT action is introduced as an explicit decision alternative. This formulation is motivated by sequential decision-making and optimal stopping theory, where delaying a decision may increase expected utility despite an associated cost [3,4]. The introduction of an explicit delay penalty is further consistent with classical stochastic dynamic programming frameworks, in which optimal policies balance immediate rewards against future expected utility. Such formulations are widely used to model time-dependent decision processes under uncertainty and provide a rigorous foundation for incorporating delay costs into autonomous decision-making models [11].

A delay penalty $D \geq 0$ is associated with the WAIT action, reflecting the operational cost of postponement. Waiting is assumed to improve the system's belief through additional observations, resulting in an updated probability p'. In a general Bayesian setting, such an update follows Bayes' rule. For the purpose of numerical exploration and interpretability, this improvement is modeled as:

$$p' = min(p + \Delta p, 1), \tag{6}$$

where Δp represents the expected information gain. This simplified representation serves as a proxy for the value of information, which has been extensively studied in decision analysis [9].

The expected utility of the WAIT action is then defined as:

$$EU_{WAIT}(p) = -D + EU_{STRIKE}(p'), \tag{7}$$

corresponding to a one-step look-ahead policy.

This formulation captures the trade-off between delay costs and the potential benefit of improved information without introducing unnecessary model complexity.

2.4 Optimal decision rule

Given the expected utilities linked with each action, the optimal decision is determined by:

$$a^*(p) = \underset{a \in \mathcal{A}}{\text{argmax}}\{EU_{STRIKE}(p), EU_{WAIT}(p), EU_{ABORT}(p)\} \tag{8}$$

This rule yields a partition of the parameter space into distinct decision regions corresponding to engagement, delay, and abort. Importantly, the explicit inclusion of the WAIT option allows for the emergence of intermediate regimes where delayed decision-making is optimal, even when the cost of waiting is non-negligible.

In practical terms, this decision rule implies that for each level of uncertainty the autonomous system evaluates all available actions and selects the one that minimizes expected loss. The resulting decision regions are not imposed heuristically, but emerge directly from utility comparisons. The inclusion of the WAIT option introduces an intermediate regime in which neither immediate engagement nor mission abort is optimal, reflecting situations where additional information can meaningfully improve decision quality. Notably, this intermediate regime persists even when waiting incurs a positive cost, illustrating that time pressure alone does not eliminate the rational value of delayed decision-making.

2.5 Numerical implementation and analysis procedure

To analyze the behavior of the proposed model, a grid-based numerical simulation is performed. The parameter space is discretized over the target legitimacy probability $p \in [0,1]$, the relative cost ratio C/B, and multiple values of the delay penalty D. For each parameter combination, expected utilities are computed and the optimal action is selected according to the decision rule above.

The numerical analysis is implemented in Python (version 3.14.2), using standard scientific computing libraries for numerical evaluation and data handling. The simulation workflow includes systematic parameter sweeping over the target legitimacy probability, relative cost ratios, and delay penalties, followed by the computation of expected utilities for each available action and the selection of the optimal decision according to the maximum expected

utility criterion. For each parameter combination, the resulting decision outcomes and associated utility values are recorded.

The simulation outputs are aggregated into a structured tabular dataset, exported in spreadsheet format, which serves as a complete numerical representation of the decision space. This table contains the parameter values, computed expected utilities, and the corresponding optimal action for each configuration, enabling reproducibility, verification, and post hoc analysis. Visualization routines subsequently process this dataset to generate decision region maps and sensitivity plots. All computations are conducted using synthetic data in order to emphasize the structural and analytical properties of the decision model rather than application-specific datasets. This implementation ensures transparent exploration of how uncertainty, cost asymmetry, and delay penalties jointly influence autonomous engagement decisions.

3. RESULTS

3.1 Baseline decision regions without delay cost

The first analysis presents the baseline decision regions obtained for a zero delay penalty ($D = 0$) and a fixed information gain $\Delta p = 0.15$. This scenario represents an idealized setting in which postponing a decision does not incur any operational cost, thereby isolating the pure effect of uncertainty and cost asymmetry on the optimal decision.

Figure 2 reveals a clear partition of the parameter space into three distinct decision regimes. For low values of the target legitimacy probability, the ABORT decision dominates across the entire range of cost ratios r. The outcome reflects a rational avoidance of engagement when uncertainty is high, regardless of the relative engagement cost. In this region, neither immediate engagement nor delayed decision-making yields a positive utility. As the probability p increases, a broad intermediate WAIT region emerges between ABORT and STRIKE.

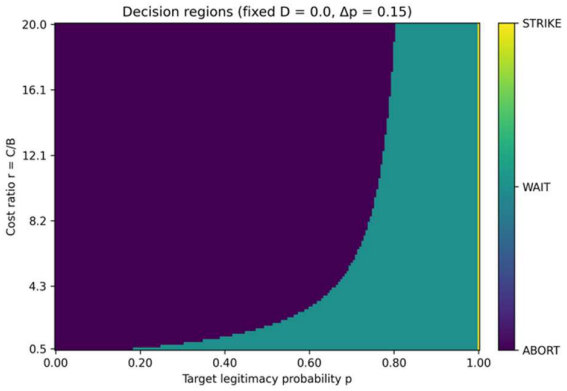


Fig. 2. Baseline decision regions (no delay cost).

This region indicates that, even when engagement is not yet justified, postponing the decision to acquire additional information is optimal. The size of this region highlights the significant value of information in environments characterized by asymmetric costs and uncertainty. In the absence of delay penalties, waiting becomes the preferred strategy over a wide range of probabilities and cost ratios.

Immediate engagement (STRIKE) is optimal only in the upper-right portion of the parameter space, corresponding to very high target legitimacy probabilities and relatively low cost ratios. This finding confirms that, even under favorable conditions, rational engagement requires a high degree of confidence when potential engagement costs are substantial.

Figure 2 establishes the fundamental structure of the decision problem and serves as a reference point for assessing the impact of increasing delay penalties in subsequent analyses. The dominance of the WAIT region in this baseline case underscores the importance of explicitly modeling delayed decision-making in autonomous systems.

3.2 Effect of non-zero delay penalties

Figures 3 and 4 illustrate the evolution of the decision regions when a non-zero delay penalty is introduced, for $D = 0.5$ and $D = 1.0$, respectively, while maintaining a fixed information gain $\Delta p = 0.15$. Compared to the baseline case discussed in Section 3.1, the introduction of a delay cost modifies the structure of the decision space but does not fundamentally alter its qualitative organization.

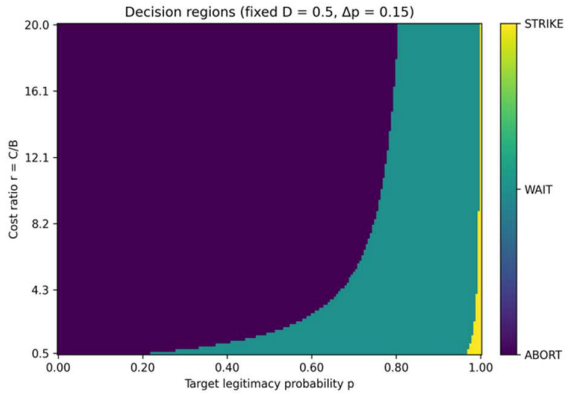


Fig. 3. Decision regions with low delay penalty.

As expected, increasing the delay penalty reduces the extent of the WAIT region. The boundary separating WAIT from STRIKE shifts toward higher values of the target legitimacy probability p, indicating that stronger evidence is required to justify postponement as waiting becomes more costly. Nevertheless, a substantial WAIT region persists across a wide range of cost ratios, particularly when the relative engagement cost is high. This demonstrates that the informational benefit associated with delaying a decision can still outweigh moderate delay costs in risk-sensitive environments.

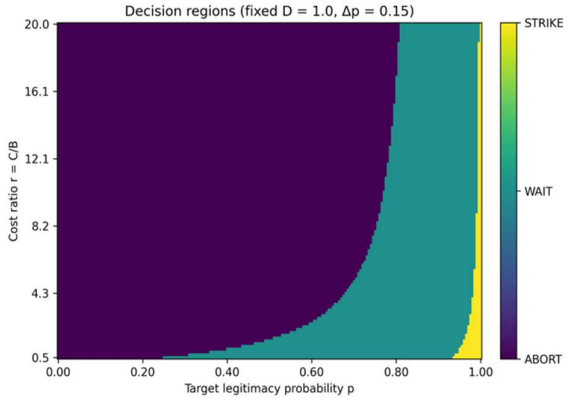


Fig. 4. Decision regions with moderate delay penalty.

The ABORT region remains unaffected by the introduction of moderate delay penalties. For low values of p, aborting the mission continues to dominate the decision space regardless of the cost ratio. This stability suggests that the decision to disengage is primarily driven by uncertainty and cost asymmetry, rather than by temporal considerations.

Immediate engagement becomes optimal only in regions characterized by both high confidence and relatively low cost ratios. Compared to the baseline case, these regions expand slightly as D increases, reflecting the growing incentive to avoid postponement. However, the gradual nature of this expansion highlights that moderate time pressure alone is insufficient to eliminate the advantage of delayed decision-making.

Thus, Figures 3 and 4 demonstrate that the proposed decision model exhibits structural robustness under moderate delay penalties. While increasing D systematically shifts decision boundaries, the presence of a clearly defined WAIT region confirms that delayed decisions remain optimal across a broad and practically relevant portion of the parameter space.

3.3 High delay cost regimes

Figures 5 and 6 illustrate the decision regions obtained for high delay penalties, corresponding to $D = 2.0$ and $D = 5.0$ respectively. These scenarios represent operational environments in which postponing a decision incurs substantial cost, such as rapidly evolving threats or limited mission windows.

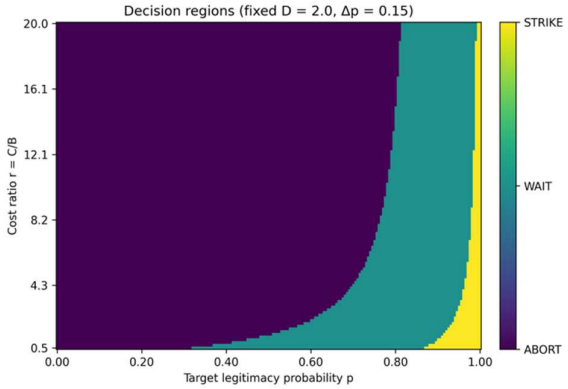


Fig. 5. High delay cost regime.

As the delay penalty increases, a pronounced contraction of the WAIT region is observed. The boundary between WAIT and STRIKE shifts markedly toward higher target legitimacy probabilities, indicating that postponement is justified only when the anticipated information gain is sufficiently valuable to offset the increased delay cost. Consequently, immediate

engagement becomes optimal over a larger portion of the parameter space, particularly for low cost ratios.

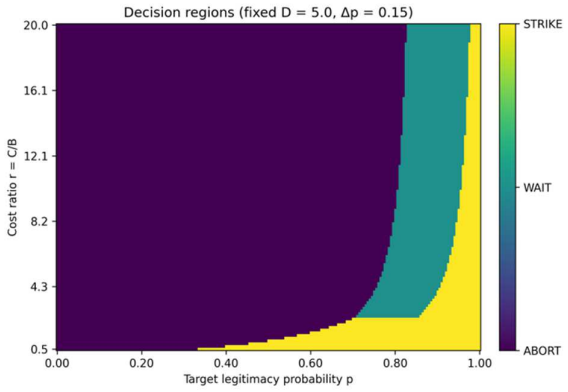


Fig. 6. Extreme delay cost regime.

Despite this trend, a key result emerges: even under severe delay penalties, the WAIT region does not vanish entirely. For intermediate probabilities and elevated cost ratios, delaying the decision remains optimal. This persistence demonstrates that the value of additional information can outweigh substantial temporal penalties when the consequences of erroneous engagement are sufficiently severe. Such behavior would not be captured by static decision rules that exclude delayed actions.

The ABORT region continues to dominate large areas of the parameter space, especially at low probabilities and high cost ratios. This dominance underscores the fundamentally conservative nature of rational decision-making in high-risk environments. Importantly, the expansion of the STRIKE region under high delay costs does not occur at the expense of ABORT, but primarily at the expense of WAIT, highlighting the distinct roles played by uncertainty, cost asymmetry, and time pressure.

Taken together, Figures 5 and 6 demonstrate that increasing delay penalties systematically alters decision boundaries without collapsing the decision problem into a purely binary choice. The continued presence of a WAIT region, even under extreme delay costs, constitutes a central finding of this study and emphasizes the necessity of explicitly modeling delayed decision-making in autonomous systems operating under uncertainty.

3.4 Sensitivity analysis with respect to the target legitimacy probability

Figure 7 provides a one-dimensional sensitivity analysis of the optimal decision as a function of the target legitimacy probability p, for a fixed cost ratio $r \approx 4.92$ and multiple values of the delay penalty D. This representation complements the two-dimensional decision maps by offering a concise view of how increasing confidence in target legitimacy influences the selected action under different temporal constraints.

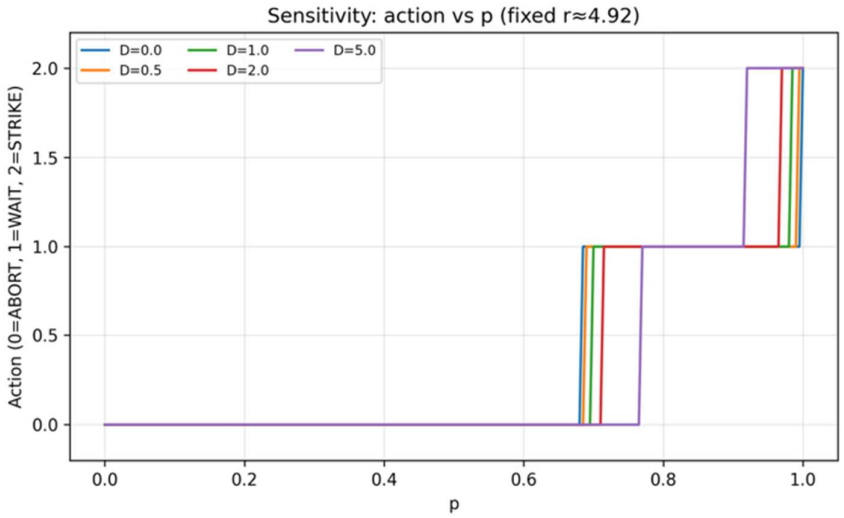


Fig. 7. Sensitivity analysis (action vs probability).

The figure exhibits a stepwise structure reflecting the discrete action space (ABORT, WAIT, STRIKE). For low values of p, the optimal decision is consistently ABORT across

all considered delay penalties. This confirms that, when uncertainty is high, neither immediate engagement nor delayed action yields sufficient expected utility, regardless of the cost of waiting.

As the probability p increases, a transition from ABORT to WAIT is observed. The location of this transition shifts systematically toward higher probability values as the delay penalty increases. This behavior indicates that postponement is justified only when the expected information gain compensates for the growing temporal cost. Nevertheless, the persistence of the WAIT region across all delay scenarios demonstrates that delayed decision-making remains a robust intermediate strategy, even under significant time pressure.

The final transition from WAIT to STRIKE occurs at very high probability levels. Notably, higher delay penalties compress the WAIT interval and cause this transition to occur earlier, reflecting the increased incentive for immediate engagement when postponement becomes expensive. However, even for the largest delay penalty considered, STRIKE is selected only under near-certainty conditions.

Overall, Figure 7 synthesizes the effects observed in the two-dimensional decision maps and confirms the monotonic influence of the delay penalty on decision conservatism. By explicitly illustrating how action selection evolves with increasing confidence, this sensitivity analysis reinforces the interpretability of the proposed model and highlights the critical role of probability thresholds in risk-aware autonomous decision-making.

4. FURTHER RESEARCH

While the proposed model provides a transparent and analytically grounded framework for autonomous engagement decisions, several directions for further research naturally emerge. First, the simplified representation of information gain may be replaced by explicit Bayesian belief updates based on sensor likelihood models, enabling a more realistic treatment of uncertainty. Such extensions would allow the integration of heterogeneous sensor reliability and context-dependent information quality into the decision process. Second, the one-step decision framework can be extended to multi-stage sequential models to account for repeated delays, cumulative delay penalties, and long-term decision strategies, offering deeper insight into temporal trade-offs in prolonged or dynamically evolving missions.

Additionally, recent research in uncertainty-aware decision-making and planning for autonomous systems highlights methods that adapt decisions dynamically based on real-time uncertainty estimates, suggesting potential hybridization with the present utility-maximization framework to enhance resilience under rapidly changing environments [12].

Future research may also investigate more complex cost structures. Rather than assuming fixed engagement benefits and costs, state-dependent or stochastic cost models could be incorporated to reflect varying collateral risk, mission priorities, or environmental conditions. These extensions would improve the representational fidelity of the model in realistic operational settings without sacrificing analytical clarity.

Additionally, alternative risk-sensitive decision criteria may be explored. In particular, coherent risk measures such as Conditional Value-at-Risk (CVaR) have been successfully applied to robust and risk-aware decision-making under uncertainty [13]. Integrating such measures into the present utility-based framework could provide additional insight into extreme-risk scenarios and tail-event behavior, while preserving analytical tractability and interpretability.

Finally, the proposed framework may be integrated with data-driven components. Learning-based methods could be employed to estimate probability distributions, cost parameters, or information gains from operational data, while retaining explicit utility maximization as the final decision mechanism. Such hybrid approaches would combine the adaptability of learning-based systems with the transparency and verifiability of decision-theoretic models.

5. CONCLUSIONS

This paper proposed a decision-theoretic framework for autonomous UAV engagement under uncertainty, explicitly accounting for asymmetric costs, probabilistic target legitimacy, and the option to delay action. Grounded in Bayesian decision theory and sequential decision-making principles, the model extends static engagement rules by treating delayed decision-making as a first-class alternative alongside immediate engagement and mission abort. This formulation enables a transparent analysis of how uncertainty, cost asymmetry, and temporal considerations jointly influence autonomous decision-making.

The numerical analysis revealed a consistent and interpretable structure of decision regions across a wide parameter space. Abort decisions dominate under high uncertainty and severe cost asymmetry, indicating that rational autonomous behavior often favors non-action in high-risk environments. Delayed decision-making emerges as an optimal intermediate strategy across many scenarios, even when postponement incurs non-negligible costs, highlighting the critical role of information value. Immediate engagement is justified mainly under conditions of high confidence and relatively low engagement costs, reinforcing the conservative nature of rational decision policies in safety-critical contexts.

A key contribution of this study lies in the explicit characterization of decision boundaries and their sensitivity to delay penalties. The results demonstrate that increasing time pressure systematically shifts engagement thresholds but does not eliminate the value of waiting when the consequences of erroneous action are severe because of the behavior that cannot be captured by static threshold-based rules.

By providing interpretable decision region maps and sensitivity analyses, the proposed approach offers a transparent alternative to black-box decision models and supports verification, validation, and human oversight. Moreover, the framework establishes a flexible foundation for future extensions, including more complex belief updates, alternative risk-sensitive criteria, and integration into real-time autonomous decision-support systems.

6. ACKNOWLEDGEMENTS

This paper was co-financed by The Bucharest University of Economic Studies (ASE) during the PhD programme.

7. REFERENCES

[1] Berger, J. O., *Statistical Decision Theory and Bayesian Analysis*, Springer-Verlag, ISBN 978-0387960982, New York, 1993.

[2] DeGroot, M. H., *Optimal Statistical Decisions*, McGraw-Hill, ISBN 978-0070169762, New York, 1970.

[3] Wald, A., *Sequential Analysis*, Courier Corporation, ISBN 978-0486780777, Mineola, NY, 2013.

[4] Bertsekas, D. P., *Dynamic Programming and Optimal Control, Volume 1*, Athena Scientific, ISBN 978-1886529434, Nashua, 2017.

[5] Pang, W., Ma, X., Liang, X., Liu, X., Dong, E., *Role-Based Bayesian Decision Framework for Autonomous Unmanned Systems*, Journal of Systems Engineering and Electronics, ISSN 1004-4132, Beijing, 2023.

[6] Vivaldini, K. C. T., Pěnička, R., Saska, M., *Decision-Making-Based Path Planning for Autonomous UAVs: A Survey*, arXiv (Cornell University), http://arxiv.org/abs/2508.09304

[7] Senthilnath, J., Harikumar, K., Suresh, S., *Metacognitive Decision Making Framework for Multi-UAV Target Search Without Communication*, IEEE Transactions on Systems, Man, and Cybernetics: Systems, ISSN 2168-2216, New York, 2024.

[8] von Neumann, J., Morgenstern, O., *Theory of Games and Economic Behavior*, Princeton University Press, ISBN 978-0691130613, Princeton, 1944.

[9] Raiffa, H., Schlaifer, R., *Applied Statistical Decision Theory*, Harvard University Press, ISBN 978-0674171555, Cambridge, MA, 1961.

[10] Howard, R. A., *Decision Analysis: Applied Decision Theory*, Proceedings of the Fourth International Conference on Operations Research, pp. 55–71, Wiley-Interscience, Boston, 1966.

[11] Ross, S. M., *Introduction to Stochastic Dynamic Programming*, Academic Press, ISBN 978-0125984201, San Diego, 2014.

[12] J. Zhou, Y. Gao, B. Olofsson, and E. Frisk, *Uncertainty-Aware Decision-Making and Planning for Autonomous Forced Merging*, arXiv (Cornell University), Oct. 2024, doi: https://doi.org/10.48550/arxiv.2410.20514.

[13] Chow, Y., Tamar, A., Mannor, S., Pavone, M., *Risk-Sensitive and Robust Decision-Making: A CVaR Optimization Approach*, arXiv (Cornell University), http://arxiv.org/abs/1506.02188

Modelare decizională bayesiană dinamică pentru angajarea autonomă a vehiculelor aeriene fără pilot în condiții de incertitudine

Abstract: Vehiculele aeriene fără pilot autonome (UAV) operează în medii caracterizate de un nivel ridicat de incertitudine, în care deciziile de angajare trebuie să echilibreze costuri asimetrice, penalizări dependente de timp și informații incomplete. Această lucrare propune un cadru decizional bayesian pentru angajarea autonomă a UAV-urilor, care integrează explicit incertitudinea privind legitimitatea țintei, costurile asociate întârzierii și opțiunea de abandon. Problema decizională este formulată cu trei acțiuni posibile: atac, așteptare și renunșare, în cadrul unei funcții de utilitate așteptată. Politicile decizionale optime sunt derivate ca funcții ale probabilității posterioare de legitimitate a țintei și ale parametrilor de cost relevanți. O analiză numerică bazată pe simulări sintetice explorează structura spațiului decizional și identifică regimuri operaționale distincte. Rezultatele arată că amânarea deciziei poate fi optimă pe intervale largi ale spațiului parametrilor, în special în condiții de incertitudine ridicată sau costuri moderate ale întârzierii, subliniind importanța unui proces decizional autonom conștient de risc și adaptiv temporal.

Emanuelle PERȚA, PhD student, Bucharest University of Economic Studies, Cybernetics and Statistics; Scientific Research Assistant, "Nicolae Bălcescu" Land Forces Academy, Department of Telecommunication, Information Technology and Cyber Security, perta.emanuelle@armyacademy.ro, +40 269 432 990, 3-5 Revoluției Street, Sibiu, Romania.

Alina TOTOI, PhD, Associate professor, "Nicolae Bălcescu" Land Forces Academy, Department of Telecommunication, Information Technology and Cyber Security, totoi.alina@armyacademy.ro, +40 269 432 990, 3-5 Revoluției Street, Sibiu, Romania.