Decision-Dominant Strategic Defense Against Lateral Movement for 5G Zero-Trust Multi-Domain Networks

Tao Li, Yunian Pan, and Quanyan Zhu

Abstract Multi-domain warfare is a military doctrine that leverages capabilities from different domains, including air, land, sea, space, and cyberspace, to create a highly interconnected battle network that is difficult for adversaries to disrupt or defeat. However, the adoption of 5G technologies in battlefields presents new vulnerabilities due to the complexity of interconnections and the diversity of software, hardware, and devices from different supply chains. Therefore, establishing a zero-trust architecture for 5G-enabled networks is crucial for continuous monitoring and fast data analytics to protect against targeted attacks. To address these challenges, we propose a proactive end-to-end security scheme that utilizes a 5G satellite-guided air-ground network. Our approach incorporates a decision-dominant learning-based method that can thwart the lateral movement of adversaries targeting critical assets on the battlefield before they can conduct reconnaissance or gain necessary access or credentials. We demonstrate the effectiveness of our game-theoretic design, which uses a meta-learning framework to enable zero-trust monitoring and decision-dominant defense against attackers in emerging multi-domain battlefield networks.

1 Introduction

The U.S. military has been undergoing a doctrine transition from traditional single to multi-domain operations or warfare (MDW), which the Army formally approved in October 2022 as its new warfighting doctrine [1]. The new doctrine defines MDW as "the combined arms employment of joint and Army capabilities to create and exploit

Tao Li (Corresponding author)
New York University, NY, 11201, e-mail: `tl2636@nyu.edu`

Yunian Pan
New York University, NY, 11201 e-mail: `yp1170@nyu.edu`

Quanyan Zhu
New York University, NY, 11201 e-mail: `qz494@nyu.edu`

relative advantages that achieve objectives, defeat enemy forces, and consolidate gains on behalf of joint force commanders," [1] which directs the service to combine and integrate air, land, sea, space, and cyberspace in all facets of operations. MDW is developed in response to the 2018 National Defense Strategy [2], shifting the previous focus of U.S. national security from addressing violent extremists worldwide to great power competition and potential conflict with near-peer adversaries across air, land, sea, space, and cyberspace.

One main impetus for this doctrine transition is the technological advances and increased complexity of modern warfare. In addition to traditional platforms such as main battle tanks and guided-missile destroyers, the rise of space, information, and artificial intelligence technologies leads to enhanced and new military capabilities, such as the Advanced Extremely High-Frequency Systems [3] powered by military satellites, the Indago quadrotor unmanned aerial systems [4], and the U.S. cyber force. By leveraging the strengths of various military capabilities across multiple domains, military forces operate through the physical dimension (air, land, sea, space), influence through the information dimension (cyberspace), and achieve victory in the human dimension.

MDW involves seamless coordination and integration of forces and assets across domains to gain a competitive advantage over adversaries. For example, ground forces may work in conjunction with air and space assets to gain situational awareness, conduct precision strikes, and provide close air support. Meanwhile, naval forces may coordinate with cyberspace capabilities to disrupt an adversary's communication networks and gain information superiority. The fifth-generation (5G) wireless technology plays an important role in MDW because it provides a network infrastructure that enables faster data transfer, greater bandwidth, lower latency, and increased capacity compared to its predecessors. With 5G networks, military units across multiple domains can access and share information in real time, creating a synergistic effect that improves situational awareness and enhances command and control. Furthermore, 5G connectivity can facilitate the communication and control of unmanned and autonomous systems powered by artificial intelligence both on the ground and in the air, enabling the integration of unmanned assets into MDW. A schematic illustration of 5G networks in MDW is presented in Figure 1

Recent years have seen the adoption and implementation of 5G networks for military applications gaining momentum. The advanced features of 5G networks, despite their contributions to coordinated MDW operations, introduce security challenges periling the efficiency and effectiveness of MDW. For example, with more devices and sensors connected to the network system, 5G networks present a larger attack surface, e.g., more potential entry points for attackers to exploit, compared to previous generations. Meanwhile, as 5G networks provide faster and more reliable connectivity, they enable more sophisticated cyberattacks, such as large-scale distributed denial-of-service attacks [5], network slicing exploitation [6], and edge computing compromise [7].

Among these cyberattacks, one critical threat is the Advanced Persistent Threat (APT). APT attacks are typically carried out by skilled and well-funded attackers who use sophisticated techniques to gain unauthorized access to sensitive informa-
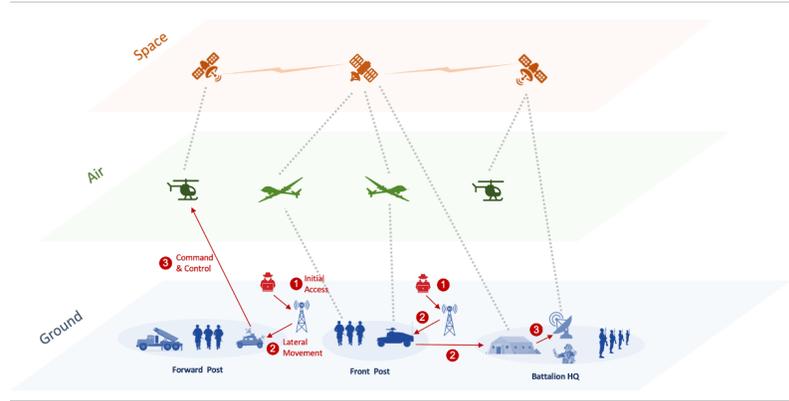
Fig. 1: An illustration of 5G Multi-Domain Networks (MDN). The army force has deployed a robust 5G communication infrastructure to facilitate seamless communication within the base and between the battalion headquarters (HQ), front lines, and forward posts. Additionally, the integration of 5G-powered satellites enables effective communication between aerial vehicles and ground forces. An APT attack can start with initial access (1), create lateral movement (2), and eventually command and control the targeted assets (3). Several paths of the attack chain are depicted, leading to the consequence of the compromise of a helicopter or misdirection to satellites

tion and systems. APT attackers may conduct extensive network reconnaissance to gather information about the 5G network and its vulnerabilities. They exploit vulnerabilities in the 5G network and gain unauthorized access to a device or system within the network to move laterally through the network and access other devices or systems within it. In 5G networks, lateral movement capabilities can be particularly dangerous, as they can allow attackers to gain access to critical systems and data within the network. For example, an attacker who gains access to a single device within a 5G network could potentially use lateral movement techniques to access other devices or systems, such as servers or databases containing sensitive or confidential data.

Since military assets and systems across various domains are connected and rely on 5G networks to exchange information and coordinate operations, the vulnerability of 5G networks can pose significant challenges in MDW. Therefore, military organizations shall prioritize the security of 5G networks in MDW and establish a proactive cyber defense in 5G networks. The primary objective of such a cyber defense is to disrupt the attacker's kill chain, which includes the following stages: reconnaissance, privilege escalation, exploitation, lateral movement, and command and control. Starting from an entry point, the attacker gains initial access to the network, conducts reconnaissance, stealthily navigates within the 5G infrastructure, and ultimately compromises the targeted asset, such as a drone or a satellite. Such adversarial behaviors are increasingly common in APTs.

To counteract the attacker's actions, the defender employs a sequence of defense actions known as the cyber defense chain, including monitoring, detection, response, and attribution. Figure 2 summarizes the kill and the defense chains. The relationship between the kill and the defense chains is competitive in nature. The kill chain aims to evade the detection from the cyber chain to reach the target, while the defense chain aims to thwart the attack before an adversary carries out the planned attack. To outmaneuver the adversary's decision-making cycle, a defender needs superior situational awareness together with fast and reliable reasoning capabilities, especially in unknown and uncertain situations to make timely and effective decisions. These desiderata are also known as decision dominance. Illustrated in Figure 2, a decision-dominant defense at the monitoring and detection stage has the capability of gathering, processing, and analyzing information from various sources to obtain a comprehensive understanding of the cyber operational environment. At the response stage, a decision-dominant defense can quickly evaluate available options, assess risks, and make informed decisions in a timely manner. As a result, it thwarts the planned attack before its execution. To achieve decision dominance, there is a need for proactive cyber mechanisms, such as cyber deception and attack engagement, to gather immediate intelligence. In addition, agility is indispensable. It allows the defender to learn, adapt, and respond to changing situations, seize opportunities, and effectively adjust strategies and tactics as required. Strategic thinking is paramount to achieving agility, involving the study of adversarial behaviors, the development of adaptive tactics, and the ability to make informed and decisive decisions.
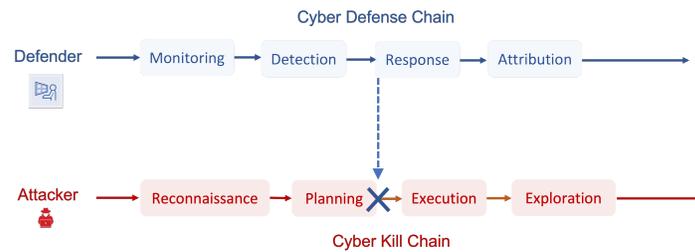
Fig. 2: An illustration of cyber kill/defense chains. The kill chain consists of crucial stages such as reconnaissance, planning, execution, and exploration. The objective of defense measures is to disrupt this kill chain by employing monitoring, detection, response, and attribution techniques. An effective defense strategy is considered decision-dominant when it efficiently acquires and processes information, enabling it to make timely decisions that outpace the attacker. For instance, the defense chain can respond swiftly to thwart the attack even before the attacker initiates the planned offensive actions

There is a pressing need for the development of a systematic approach to establishing decision-dominant mechanisms for the defense of 5G networks. Game theory offers a promising solution in this regard. Not only does game theory nat-

urally provide a framework for designing tactics in competitive environments, but recent advancements in dynamic games, learning, and their intersection with modern machine learning techniques enhance the reasoning capabilities of agents. This enables a formal and agile approach to achieving rapid decision-making. For instance, recent studies [8, 9, 10, 11] have introduced a class of dynamic games that effectively capture the evolving interactions between defense and kill chains. The concept of non-equilibrium has been proposed to derive solution concepts based on players' behaviors. This concept holds significant implications for cybersecurity applications, particularly when the interactions between attackers and defenders may be limited and indirect.

Another significant advantage of utilizing game-theoretic models is their strong epistemic foundation, which allows for explicit modeling and analysis of scenarios involving information asymmetry and the pace of decision-making. These models find wide applicability in 5G security networks. Information asymmetry arises from the fact that neither party possesses a comprehensive view of the entire 5G network. Instead, each party gathers partial observations through reconnaissance (the attacker) or monitoring (the defender). To effectively outmaneuver the adversary, the defender must establish an information advantage by actively acquiring information during the monitoring process. This proactive approach enables the defender to gain high-confidence situational awareness of the network system and adversarial behaviors. However, it is important to note that having an information advantage alone does not necessarily guarantee the defender an upper hand in cyber defense. Another crucial aspect that holds equal importance is the pace of decision-making. The defender faces a disadvantage if the attacker manages to execute the attack successfully before an adequate response can be mounted. In this regard, game theory frameworks provide a means to comprehensively capture the end-to-end decision-making process, encompassing information acquisition, learning, and decision-making. It provides a theoretical underpinning for understanding the fundamental tradeoff among these factors and a holistic approach to modeling and devising tactics across all stages.

One implicit assumption underlying the defense against APTs is that the attacker possesses the necessary capabilities to acquire initial access and credentials, and then establish a foothold within the network. We cannot stop the attack from getting into the network. This assumption forms the basis of the zero-trust security doctrine, which emphasizes the need to trust no entity by default and requires organizations to verify and authenticate all users, devices, and activities, regardless of their location or origin. Recognizing the importance of assuming a reasonable capability of adversaries in developing effective defenses, the concept of zero-trust doctrine can also be integrated into game models by establishing relevant adversarial models. By incorporating the principles of zero trust, game models can create decision-dominant zero-trust policies to defend against APTs in 5G networks.

To this end, we propose a decision-dominant zero-trust defense (DD-ZTD) against adversarial attacks in 5G networks in MDW to strike the right balance between information acquisition and fast decision-making. DD-ZTD is built on a game-theoretic framework that captures the information asymmetry and the competitive nature of cyber defense. Following the "never trust, always verify" principle [12], zero-trust

defense (ZTD) equips the defender with a proactive information processing mechanism when operating with incomplete information about the attacker's intentions, capabilities, and actions, which is crucial to develop strategies that account for the information asymmetry.

The ZTD problem of the 5G network is modeled as an *asymmetric information Markov game* (AIMG) between the defender and the attacker. Thanks to its great expressivity, AIMG offers a comprehensive characterization of various information structures in cyber defense, which facilitates defense design in various security contexts. Furthermore, the equilibrium notion in AIMG lays a theoretical underpinning of an adaptive ZTD in the presence of information asymmetry. Powered by recent advancements in machine learning, the proposed game-theoretic ZTD framework exhibits great potential in devising a generalizable intelligent defense against a wide range of cyber attacks arising from a variety of network systems possibly unknown to the defender beforehand.

To outpace the attacker in the cyber kill chain, ZTD is further augmented by decision dominance (DD), where DD accelerates the defense decision-making in ZTD. As its name suggests, DD makes the defender the dominant player in the dynamic game by taking decisive actions based on acquired partial information with high confidence before the attacker compromises the network system, sharing the same spirit of the motto "first look, first shot, first kill."[13]. Such strategic dominance is achieved by game-theoretic calculations where the defender takes into account the attacker's decision-making process. DD amounts to an optimal stopping (Dynkin's) game problem, which essentially captures the defender's strategic anticipation of the opponent's stopping criterion, as well as the fundamental tradeoff between the benefits and harm of lingering in the interaction, which is ubiquitous in the cyber security domain. The equilibrium notion for DD enables the defender to make opponent-independent stopping decisions based on the payoff evaluation for the underlying cyber kill chain process while making the monitoring and investigation as effective as possible.

The rest of this chapter is organized as follows. Section 2 provides an overview of multi-domain warfare and associated 5G networks across multiple domains, laying the context for further discussions. Section 3 articulates the emerging security challenges in 5G networks, particularly the advanced persistent threats (APT). To address these security issues, we propose a decision-dominant zero-trust defense for 5G networks in Section 4, where the game-theoretic conceptualization is presented. Section 5 and Section 6 dive into the details of the zero-trust defense and the decision-dominance concept in detail, respectively, where case studies of the proposed DD-ZTD are presented.

2 Multi-Domain Warfare and 5G Networks

This section briefly overviews multi-domain warfare and the associated 5G communication networks.

2.1 Multi-domain Warfare

Multi-domain warfare (MDW), a new operation concept designated by the U.S. Army [14], refers to the combined arms employment of military capabilities straddling multiple domains to create and exploit a decisive advantage over an adversary. Unlike traditional warfare, where operations are conducted within a single domain, MDW rests on synthesizing various military capabilities across five warfighting domains: land, sea, air, space, and cyberspace.

The backbone of MDW is the coordination and integration among different military units from multiple domains, leading to joint operations where various military services, such as the army, navy, air force, and space force, work together collaboratively. By operating across multiple domains, military forces can disrupt an adversary's operations and degrade their ability to fight.

2.2 5G Multi-Domain Networks

One challenge to achieving real-time coordination and integration in multi-domain warfare is the lack of network infrastructure to support interoperability among military units using different communication systems, making coordinating actions across multiple domains difficult. The fifth generation (5G) wireless communication technology plays a vital role in multi-domain warfare. It provides a network infrastructure that enables faster data transfer speeds, greater bandwidth, lower latency, and increased capacity and reliability than previous generations of mobile networks. Thanks to its advanced features, 5G technology provides the foundation for faster, more connected, and more capable military operations across multiple domains, leading to improved situational awareness, enhanced command and control, precise targeting, integration of unmanned systems, and support for emerging technologies like the internet of battlefield things(IoBT). We elaborate on these aspects in the ensuing paragraphs. Figure 1 presents a schematic illustration.

Situational Awareness 5G MDN can support the transmission of large volumes of data in real time. This enables the rapid exchange of information between sensors, platforms, and command centers across different domains. Improved situational awareness allows military commanders to make more informed decisions and respond promptly to changing battlefield conditions.

Precise Targeting The low latency and high bandwidth of 5G networks enable the real-time transmission of sensor data and imagery, supporting the precise targeting of enemy assets. This enhances the effectiveness of kinetic operations, such as precision strikes, and improves the accuracy of intelligence, surveillance, and reconnaissance (ISR) capabilities.

Command and Control 5G networks can facilitate seamless communication and coordination between military units and commanders across domains. Reliable and

low-latency connectivity enables the transmission of commands, orders, and mission-critical data, enhancing command and control capabilities in multi-domain operations.

Integration of Unmanned Systems and IoBT 5G connectivity can facilitate the communication and control of unmanned systems and autonomous vehicles, both on the ground and in the air. This enables the integration of unmanned assets into multi-domain operations, enhancing their situational awareness, coordination, and responsiveness. In addition, 5G connections among a massive number of devices and sensors can be leveraged to create a comprehensive network of interconnected assets. This integration allows for better monitoring, management, and control of unmanned systems, autonomous vehicles, and other IoT devices across domains.

3 Emerging Security Challenges in 5G Multi-Domain Networks

5G networks represent a significant advancement in technology, offering functionalities that set them apart from previous generations. In the context of multi-domain warfare, it is crucial to examine the vulnerabilities inherent in 5G networks, as they can be exploited to form an APT kill chain. This section will delve into the vulnerabilities stemming from APIs, network slicing, and the supply chain.

3.1 Security of 5G Multi-Domain Networks

5G networks play an important role in MDW as they provide a network infrastructure that enables faster communication, greater bandwidth, and lower latency between different military units compared to previous generations of mobile networks. With 5G technology, military personnel can access and share information in real-time, allowing for faster decision-making and more efficient deployment of resources. For example, a military unit is conducting a mission in an urban environment that involves ground troops, drones, and surveillance equipment. The troops on the ground need to communicate with each other in real time while also receiving information from the drones and surveillance equipment to coordinate their actions.

Moreover, 5G technology allows for the use of advanced technologies such as drones, autonomous vehicles, and augmented reality, which can be used to gather intelligence, conduct surveillance, and engage in combat operations. These technologies rely on high-speed, low-latency networks to function effectively, and 5G provides the necessary infrastructure to support their deployment. For example, during the U.S. military's operations in Afghanistan, the 5G-satellite communication network was used to provide real-time communication and intelligence sharing between ground forces, aircraft, and command centers. The system enabled military

forces to coordinate their actions across different domains while also providing them with the information and intelligence needed to make informed decisions.

In addition to its communication capabilities, 5G-supported satellite networks also have the ability to support other mission-critical functions, such as intelligence gathering and surveillance. The system's high-capacity communication services and advanced technology make it a critical enabler for multi-domain warfare, providing military forces with the network infrastructure needed to support real-time communication and information sharing across different domains.

Recent years have seen that the adoption and implementation of 5G networks for military applications are gaining momentum. As military forces become more reliant on 5G networks, they also become more vulnerable to cyber-attacks. To achieve multi-domain warfare, military forces need to develop robust cybersecurity measures to protect their 5G networks and systems from cyber threats. One critical threat is APT attacks on 5G networks. APT attacks are typically carried out by skilled and well-funded attackers who use sophisticated techniques to gain unauthorized access to sensitive information and systems. APT attackers may conduct extensive network reconnaissance to gather information about the 5G network and its vulnerabilities. They exploit vulnerabilities in the 5G network and gain unauthorized access to a device or system within the network to move laterally through the network and access other devices or systems within it. In 5G networks, lateral movement capabilities can be particularly dangerous, as they can allow attackers to gain access to critical systems and data within the network. For example, an attacker who gains access to a single device within a 5G network could potentially use lateral movement techniques to gain access to other devices or systems, such as servers or databases, which contain sensitive or confidential data.

3.2 5G Threat Landscape: Vulnerabilities and Kill Chain

The emergence of 5G technology represents a significant departure from previous mobile generations, bringing with it a distinct set of security requirements. This is particularly crucial for military users who often necessitate tailored and specialized services to address their unique operational needs. There are several key threats associated with 5G networks beyond general cybersecurity threats (e.g., unauthorized access, human errors, and misconfigurations). Various threat frameworks are available to aid in analyzing these threats, such as those provided by MITRE Fight and 3GPP's Security Assurance Specifications (SCAS) and Technical Specification (TS) 33.501.

One prominent threat to 5G networks is virtualization threats, which impact virtual machine (VM) and container service platforms, affecting various aspects of 5G, including the Core, RAN, MEC, Network Slicing, Virtualization, and Orchestration and Management. These threats encompass DoS attacks, VM/container escape, side-channel attacks, and misconfigurations by cloud service consumers. For instance, extreme resource consumption by one tenant in a multi-tenant virtualization

environment can lead to a DoS event for neighboring tenant systems, impeding mission functionality. Similarly, colocation attacks, such as VM/container escape or side-channel attacks, can compromise neighboring compute workloads, resulting in resource deprivation, lateral movement, and compromising data confidentiality, integrity, or availability. A side-channel attack on 5G RAN or Core functions could allow bypassing user account permissions, virtualization boundaries, or protected memory regions, thereby exposing sensitive information.

One type of threats is on 5G network slices. These threats may exploit weaknesses in the network slice's configuration, protocols, or applications, potentially leading to unauthorized access, data breaches, or service disruptions within that particular slice. To combat this threat, slice isolation is a promising approach. It involves creating and maintaining separate virtual network slices within the 5G infrastructure. By isolating slices, potential interference or vulnerabilities in one slice are contained, ensuring the integrity and security of other slices.

As 5G networks utilize application programming interfaces (APIs) for communication and interaction between different components, several potential threats can arise. These include DoS attacks targeting 5G APIs by overloading them with a high volume of requests or exploiting API vulnerabilities to exhaust system resources. Attackers can also exploit API vulnerabilities by abusing or misusing them to gain unauthorized access, manipulate data, or disrupt services. This can involve sending malicious API requests, performing injection attacks, or overwhelming the API with excessive requests (API flooding).

The increasing complexity of 5G networks involves a vast ecosystem of suppliers and vendors. Security vulnerabilities in the supply chain can lead to compromised components or malicious software being introduced into the network infrastructure, posing significant risks. For example, the presence of counterfeit or substandard components in the 5G supply chain poses significant risks to network security and integrity. These components may not meet the required quality standards or security specifications, making them susceptible to exploitation and compromise. Unauthorized actors could exploit these vulnerabilities to gain unauthorized access or control over the network infrastructure, potentially leading to data breaches, service disruptions, or unauthorized surveillance.

In addition to counterfeit components, there is a risk of introducing malicious software or hardware into the 5G supply chain. This can occur through intentional modifications or the inclusion of backdoors that provide unauthorized access points. Threat actors can exploit these vulnerabilities to infiltrate the network infrastructure, compromise the confidentiality, integrity, and availability of data, or gain unauthorized control over critical network functions.

Supply chain security risks can also originate from third-party providers involved in the network deployment, such as installation contractors or maintenance service providers. Inadequate security measures implemented by these third parties, insider threats, or the compromise of their systems can introduce vulnerabilities into the 5G network. Weaknesses in the security practices of these entities can be exploited by threat actors, compromising the overall security of the network.
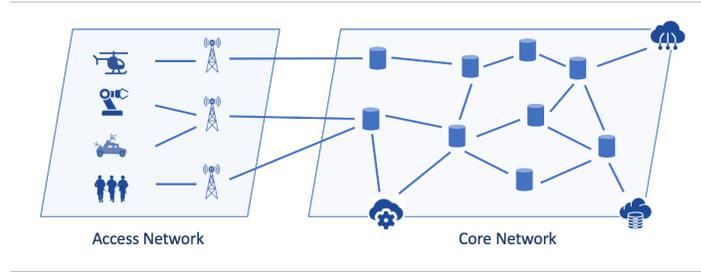
Fig. 3: An illustration of 5G network consisting of access network and core network as two major components. The core network has the functionalities of control plane and user plane separation, network functions virtualization (NFV), network slicing, mobility management, and multi-access Edge Computing (MEC)

The combination of vulnerabilities in API, supply chain, and network slicing, along with others, can be exploited by an Advanced Persistent Threat (APT) attack to form a comprehensive kill chain. Fig. 3 provides a visual representation of a baseline 5G network, where UEs utilizing O-RAN technology connect to the 5G core networks. This interconnected infrastructure presents an attack surface that an adversary can leverage to target specific entities. By capitalizing on the identified vulnerabilities, an attacker can exploit weaknesses in the API layer, infiltrate compromised components introduced through the supply chain, and exploit insufficient isolation or monitoring within the network slicing architecture. This enables the attacker to establish a persistent presence within the network and navigate through various stages of the kill chain to reach their intended target. Fig. 1 has illustrated the potential attack path an adversary may take, highlighting the entry points, lateral movement, and potential impact on the 5G network. Understanding and visualizing this attack surface assists in identifying critical areas for security enhancements and mitigations.

Zero-trust policies can be implemented to counteract such threats. It aims to establish clear rules and guidelines for access, authentication, and data protection within the network. These policies define which individuals or entities have access to specific resources, under what conditions, and the level of authorization required. It is crucial for the policy to align with the organization's security objectives and regulatory requirements. Regular monitoring of network traffic, user behavior, and access logs is essential to promptly identify any anomalies or potential security breaches. Additionally, it is important to periodically review and update the Zero Trust policy to adapt to evolving threats and changes in the network environment.

4 Decision-Dominant Zero-Trust Defense: A Game-Theoretic Framework

This section presents a high-level overview of the proposed decision-dominant zero-trust defense (DD-ZTD) in 5G multi-domain networks, arguing that the proposed game-theoretic framework leads to a unified framework for cyber defense in 5G networks.

4.1 Decision Dominance

Decision dominance refers to the ability of a defender to outmaneuver the adversary's decision-making cycle by possessing superior situational awareness and efficient reasoning capabilities. It involves making timely and effective decisions, particularly in unknown and uncertain situations, in order to gain an advantage over the attacker. To achieve decision dominance, a defense strategy needs to excel in two stages: monitoring and detection and response. In the monitoring and detection stage, a decision-dominant defense can gather, process, and analyze information from various sources to obtain a comprehensive understanding of the cyber operational environment. This enables the defender to proactively identify and assess potential threats. In the response stage, a decision-dominant defense can swiftly evaluate available options, assess risks, and make informed decisions in a timely manner. By doing so, it can effectively thwart planned attacks before they are executed. Achieving decision dominance requires proactive cyber mechanisms like cyber deception and attack engagement to gather immediate intelligence. Agility is also crucial, allowing the defender to learn, adapt, and respond to changing situations, seize opportunities, and adjust strategies and tactics as necessary.

Zero-trust decision-dominance strategies refer to a specific type of decision-dominance strategy that operates on the assumption of the presence of adversaries at all times. These strategies are particularly critical for securing 5G networks, given the expanding attack surface and the significant number of IoT devices deployed in battlefield environments. Implementing these strategies requires strategic thinking and continuous monitoring of device behaviors to assess their trustworthiness. Timely evaluation and rapid response capabilities are essential in terms of network configuration and access control policies to counteract adversaries before they can execute their planned attacks. To ensure effective implementation, it is necessary to establish quantitative and formal frameworks that incorporate zero-trust decision-dominance into 5G network security policies. These frameworks provide a structured approach to design and enforce robust security measures that align with the principles of zero trust, enhancing the overall resilience and protection of 5G networks in dynamic threat environments.

4.2 Conceptualization of Decision-Dominant Zero-Trust Defense

One of the primary objectives of this book chapter is to develop a quantitative framework that formalizes the decision-making process for zero-trust defense. The inherent competition between attackers and defenders naturally gives rise to a dynamic game environment that reflects the win-lose nature of multi-stage interactions. To account for the information asymmetry between the players resulting from differences in monitoring and sensing capabilities, we propose a dynamic game of asymmetric information. In this game, players utilize the information available to them through the established information structure to infer unknowns. Variations in the information structure lead to differing belief structures. Players make decisions based on their beliefs, resulting in new observations in subsequent rounds of interaction and the formation of updated beliefs. It is evident that there exists interdependence between the beliefs and actions arising from the players' chosen strategies. The solution concept for the game necessitates consistency between the agents' beliefs and their optimal effort strategies. This concept gives rise to the notion of Bayesian Nash equilibrium, which serves as the foundation for developing algorithms to implement game-theoretic solutions in practical scenarios.

It is important to note that belief formation stems from incomplete information regarding the other agent. In our case, the incomplete information pertains to the behavior of the other player. Thus, it can also be seen as a process of establishing trust in the other player. This naturally aligns with the concept of zero trust, which requires the defender to distrust users or third-party players in the network despite their credentials. At the outset, the true identity must be considered unknown and untrusted, and the evaluation of a player's trustworthiness epitomizes the principle of zero trust. The baseline equilibrium concept is established using Bayesian rationality, where Bayes' law is employed to update beliefs whenever new observations are obtained by the players. In practice, this baseline can be replaced with a machine-learning approach for inference. In modern scenarios, vast amounts of data are collected from numerous users interacting with the system. These data can be incorporated into game-theoretic models, facilitating the practical application of equilibrium solution concepts. Detailed models and their applications to lateral movements will be discussed in the subsequent section.

In order to accommodate the requirement of quick decision-making in decision-dominant scenarios, the game becomes dynamic and no longer has a fixed horizon. In this type of game, known as a stopping time game, players have the ability to choose when to cease observations and make their decision. The advantage of stopping early lies in determining the payoffs, but there is a risk of uncertainties that may lead to higher payoffs if the decision is postponed. However, it is important to note that the other player also has the capability to terminate the game. If the attacker terminates the game prematurely, the defender would be in a passive position. Thus, the competitive nature of the game naturally leads to a decision-dominant scenario. The defender's reasoning involves inferring the opponent's strategies based on the observations and, in the meantime, trades off between the probable stopping by the attacker as well as the low payoff as a result of early stopping. To formally capture

this dynamic, we introduce a stopping-time game in the ensuing section, with the aim of creating decision-dominant strategies. The associated Nash equilibrium solution concept allows us to reason formally about the active and passive situations of the defender, referred to as defender dominance and adversary dominance, respectively. The baseline analysis provides insights into the necessary structures for developing winning solutions, including the payoff structures, information structures, and inference mechanisms. This analysis also establishes a theoretical foundation for understanding the fundamental limits of strategic decision dominance in the face of a strategic adversary. By integrating decision-dominant strategies with zero-trust defense strategies within the baseline framework, we can establish a symbiotic relationship between the two. Additionally, the consolidation and integration of data analytics can pave the way for the development of practical algorithms in the future.

The proposed framework in this book chapter is solidly built on the recent development of game-theoretic models for cybersecurity. Recent advances have witnessed the growth in their application to assess security risks, design protection mechanisms, and inform policy making for communication networks [15, 16, 17], Internet of things [18, 19, 20], power and energy systems [21, 22, 23, 24], manufacturing and robotics [25, 26, 27, 28], supply chains [29, 30, 31], and transportation networks [32, 33, 34]. Game theory has also provided theoretical foundations for cyber deception [35, 36, 37, 38], moving target defense [39, 40], and human behaviors [5, 41, 42]. Both decision-dominance and zero-trust defense possess distinct characteristics that necessitate specific game structures to capture their essential features and provide valuable insights. In this context, our focus lies on two types of game structures: the game of asymmetric information and stopping time games. This chapter not only applies these game structures to 5G zero-trust security problems but also contributes to a novel class of game-theoretic frameworks, pushing the boundaries of game theory forward.

Our contribution primarily revolves around the creation and analysis of stopping-time games within the framework of asymmetric information dynamic games. By incorporating asymmetric information into these games, we introduce a new dimension that enhances our understanding of strategic interactions. Furthermore, we consolidate the fields of meta-learning and explainable learning within the domain of asymmetric information games, fostering a comprehensive approach to game analysis. Through these contributions, we aim to extend the frontiers of game theory, providing researchers and practitioners with valuable tools to tackle decision-dominance and zero-trust defense challenges effectively.

5 Zero-Trust Defense

With a growing threat landscape and attack surfaces in 5G networks, traditional perimeter-based defense, a static defense mechanism, has become inadequate in the face of sophisticated cyber attacks, such as APTs. Advanced attackers can evade traditional intrusion detection at the perimeter, obtain privileges as an insider with

stolen credentials, and move laterally within the network. In response to the vulnerabilities in the static defense, zero trust emerges as a promising security framework, assuming that no entities can be trusted and therefore requiring verification processes for every incoming access request [12].

Zero-trust defense (ZTD) consists of two components: trust evaluation and access policy. Square one of ZTD is to quantitatively establish the trustworthiness of each entity in the network, which is highly nontrivial in 5G networks with large-scale heterogeneous network entities. Due to the increasing network connectivity, the defender can only acquire limited partial observations of the user's trace through methods such as Intrusion Detection Systems [43], and Security Information and Event Management [44]. These limited observations create **information asymmetry**, complicating the defender's decision-making, and a quantitative metric measuring the user's trustworthiness using partial observations is indispensable.

With the trust evaluation, the defender can enforce different policies for access to network resources. What distinguishes ZTD from the perimeter-based one is that the trust evaluation and the access policy, together with the network monitoring unit, constitute a feedback loop shown in Figure 4. As new observations are fed into the evaluation unit, the defender adjusts the trust and the access policy accordingly, leading to a dynamic defense. This section articulates a game-theoretic framework (see Definition 1) for ZTD design in 5G networks, which offers a natural set of tools to capture the information asymmetry and the competitive nature of the two parties in dynamic environments.
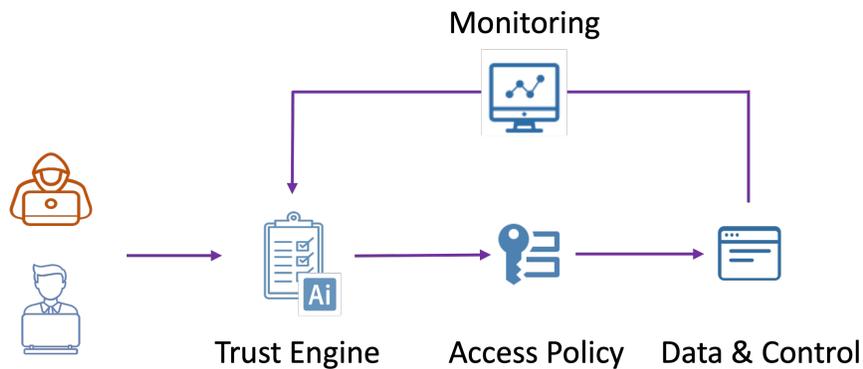
Fig. 4: An illustration of the feedback loop in zero-trust defense (ZTD) architecture. Unlike the perimeter-based defense, ZTD dynamically evaluates the trustworthiness of the user using feedback from the security monitoring system, such as SIEM [44]. Based on the trust evaluation, the access policy either grant or deny access

The proposed game-theoretic framework provides a theoretical underpinning of adaptive and strategic ZTD built upon the notion of perfect Bayesian Nash equilib-

rium (see Definition 7) in the face of asymmetric information. This equilibrium-based ZTD can be further augmented with modern machine-learning (ML) methodologies providing an end-to-end automated network defense (see Section 5.3), generalizing to adversarial scenarios unseen in the pre-training stage. As advanced ML machinery enters the picture, the ZTD architecture grows opaque to human operators. To make ML-based ZTD itself trustworthy to humans, it is necessary to increase the explainability and accountability of learning-based ZTD, which is discussed at the end of this section.

5.1 Information Asymmetry in Zero-Trust Defense

As a prevailing phenomenon in security applications [45], information asymmetry refers to the fact that one party is better informed than the other party at the point of decision-making. To facilitate our discussion, we use the notion **information structure** [45] to capture the player's observations and knowledge throughout the decision-making process, which is mathematically a set of random variables whose realizations can be observed by the player [45]. We first present a bird's eye view of asymmetric information structures in the cyber defense of 5G networks, and mathematical definitions and arguments are deferred to Definition 1 and the ensuing remarks.

Compared to its predecessors, 5G networks enjoy increasing capacity and reliability that can support a massive number of heterogeneous devices. Consequently, it becomes prohibitive, if not impossible, for either the defender or the attacker to acquire a holistic view of the underlying network. The resulting information structures of both parties' partial observations display complexities to various extents, which can be categorized according to different taxonomies. We here present two taxonomies based on the notion of information superiority proposed in [45]: one player is said to be informationally superior to the other if its information structure is a superset of its counterpart.

Depending on which party acquires the information superiority, information asymmetry includes **one-sided** and **double-sided** information asymmetry. One-sided information asymmetry refers to a situation where one party achieves information superiority over the other. If no one is informationally superior, then the resulting situation is of double-sided information asymmetry, where both parties acquire private information hidden from the other [46].

Depending on whether the information superiority is rooted in the knowledge or the observation, information structures can be categorized into **incomplete** and **imperfect** information structures. Knowledge is endogenous, reflecting the player's comprehension of the decision-making process. The incomplete information points to the player's uncertainty regarding the other's decision-making capabilities and incentives. In contrast, observation is exogenous, referring to the player's awareness of events that have previously occurred. Imperfect information refers to the situation where the player is unaware of some events in the decision-making.

As one shall see later in the running example in Section 5.2, the aforementioned information structures are prevalent in network defense. To systematically investigate information asymmetry in the cyber defense of 5G networks, we propose the asymmetric information dynamic games in the following, laying a mathematical foundation to facilitate ZTD design under sophisticated information structures, which is visualized in Figure 5.
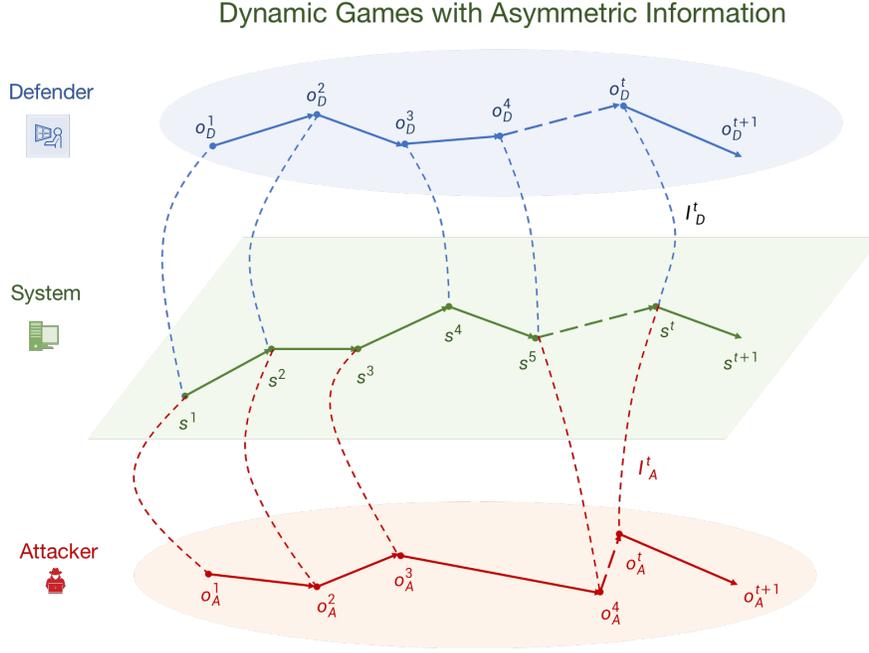
Fig. 5: An illustration of asymmetric information dynamic games defined in Definition 1. Let D and A denote the defender and the attacker, respectively. Under asymmetric information structures I_D^t and I_A^t, the two players have disparate partial observations, denoted by o_D^t, o_A^t on the system operation s^t. In zero-trust defense, the defender must infer the attacker's intention, assign trust scores, and determine the access policy based on its limited observations, which calls for efficient and adaptive trust evaluation and policy learning

Definition 1 (Asymmetric-Information Markov Game). An asymmetric-information Markov game (AIMG) $\mathcal{G}$ is given by the following tuple

$$\mathcal{G} := \langle \mathcal{N}, \Omega, \rho, \mathcal{S}, (O_i)_{i \in \mathcal{N}}, (\mathcal{A}_i)_{i \in \mathcal{N}}, P, (u_i)_{i \in \mathcal{N}}, (\sigma_i)_{i \in \mathcal{N}}, (I_i)_{i \in \mathcal{N}}, H \rangle,$$

where the definition of each component within the tuple is as below. It is assumed every set is discrete and finite. Let $t \in \mathbb{N}_+$ be the time index.

- $\mathcal{N} = \{D, A\}$ is the decision-maker (player) set, including the defender and the attacker, denoted by D and A, respectively. For simplicity, we consider a single attacker within the network, and the generalization to the case where multiple attackers coexist is straightforward.
- Ω is the attacker's type space, and its typical element ω indicates its attack capability (e.g., stealthiness) and objective (e.g., data breach). To simplify the exposition, the normal user is also treated as one type of attacker without malicious intentions or attack capabilities.
- ρ is the type distribution over Ω, and $\rho(\omega)$ implies the probability of a certain attacker ω appearing in the network.
- $\mathcal{S}$ denotes the state space with its typical element s representing the operation status of the network.
- O_i denotes the observation space, and its typical element o_i represents the player i's partial observation.
- $\mathcal{A}_i$ is the action space of the player i.
- $P : \mathcal{S} \times (\mathcal{A}_i)_{i \in \mathcal{N}} \times \Omega \to \Delta(\mathcal{S})$ is the state transition function, depicting how the network operation evolves under the joint force of the defense and attack. To be specific, $P(s^{t+1}|s^t, a_D^t, a_A^t, \omega)$ gives the probability that s^{t+1} emerges after the two players execute a_D^t and a_A^t at the state s^t.
- $u_i : \mathcal{S} \times (\mathcal{A}_i)_{i \in \mathcal{N}} \times \Omega \to \mathbb{R}$ is the instantaneous cost of the player i.
- $\sigma_i : \mathcal{S} \times (\mathcal{A}_i)_{i \in \mathcal{N}} \times \Omega \to \Delta(O_i)$ is the observation function, and $\sigma_i(o_i^t|s^t, a_D^t, a_A^t, \omega)$ denotes the probability of observing o_i^t when the underlying state is s^t.
- I_i is a set-valued mapping, characterizing the information structure of the player i throughout the Markov game. Let $\mathcal{H}^t := \{\omega, [s^k(a_i^k o_i^k)_{i \in \mathcal{N}}]_{k=1}^{t-1} s^t\}$ be the history of the gameplay up to time t, then $\mathcal{I}_i^t := I_i(\mathcal{H}^t) \subset \mathcal{H}^t$ presents the player's partial observation of the play.
- H is a constant, denoting the horizon length of the game, i.e., the operating lifetime of the network.

The AIMG unfolds as follows. In the first stage, a type-ω attacker is realized according to the distribution ρ, and the network state s^1 is initialized. At the time t, each player implements an action a_i^t from the action set $\mathcal{A}_i$ based on the information structure $\mathcal{I}_i^t$. Then, the state evolves to s^{t+1}. This procedure repeats until the game reaches the end of the horizon. The goal of type-ω attacker is to find a policy $\pi_A : \mathcal{I}_A^t \to \Delta(\mathcal{A}_A)$ within a specified policy class Π_A such that the cumulative cost is minimized:

$$\min_{\pi_A \in \Pi_A} \mathbb{E}\left[\sum_{t=1}^{H} u_A(s^t, a_D^t, a_A^t, \omega)\right], \tag{1}$$

where the expectation is taken over Borel probability measures in AIMG, including the transition P, the observation functions $(\sigma_i)_{i \in \mathcal{N}}$, and the policies $(\pi_i)_{i \in \mathcal{N}}$.

The defender's objective is more involved than (1) due to the lack of information on the attack type, and a generic characterization is given by (2), where the notations are in a similar vein of (1), except that the inner expectation $\mathbb{E}_{\omega \sim \mathcal{T}}(\cdot)$ is taken over the hidden type ω with respect to the defender's subjective belief $b^t \in \Delta(\Omega)$ based on the observations $\mathcal{I}_D^t$. Such a belief constitutes the defender's trust evaluation of the user, and a mathematical characterization is presented in Definition 2.

$$\min_{\pi_D \in \Pi_D} \mathbb{E} \left\{ \sum_{t=1}^{H} \mathbb{E}_{\omega \sim b^t} \left[u_D(s^t, a_D^t, a_A^t, \omega) \right] \right\}. \tag{2}$$

Definition 2 (Trust and Trust Engine). The trustworthiness of the user at time t is defined as a probability measure over the type space $b^t \in \Delta(\Omega)$, which is determined by the defender's trust engine Φ that maps the information structure $\mathcal{I}_i^t$ to the trustworthiness $b^t = \Phi(\mathcal{I}_i^t)$. The set of beliefs $\{b^t\}_{t=1}^{H} \in \Delta(\Omega)^H$ is referred to as the trust evaluation.

The trust metric b we consider is a probability measure, and $b(\omega)$ depicts the defender's subjective belief over the hidden type ω, also referred to as the trust score [47]. With the trust evaluation, the defender can determine the access policy $\pi_D(\mathcal{I}_i^t, b^t)$ based on its observation, which, together with the trust engine, constitutes a zero-trust defense mechanism. A mathematical definition is given below.

Definition 3 (Zero-Trust Defense). The zero-trust defense is defined as a pair of the trust engine $\Phi : \cup_{t=1}^{H} \{\mathcal{I}_i^t\} \rightarrow \Delta(\Omega)$ and the access policy $\pi_D : \cup_{t=1}^{H} \{\mathcal{I}_i^t\} \times \Delta(\Omega) \rightarrow \Delta(\mathcal{A}_D)$.

Before elaborating on the two critical components of ZTD in Subsection 5.3, we first remark on the expressive power of AIMG in modeling the cyber defense of 5G networks under complex information structures. In particular, Definition 1 leads to a systematic characterization of various information structures, such as one/double-sided information asymmetry and incomplete/imperfect information.

Definition 4 (One/Double-sided Information Asymmetry).
The player i is said to be informationally superior than j if $\mathcal{I}_j^t \subsetneq \mathcal{I}_i^t$, for all t. This information asymmetry is one-sided since the player i is always

better informed than its opponent. If there exists t such that $\mathcal{I}_i^t \setminus \mathcal{I}_j^t \neq \varnothing$ and $\mathcal{I}_j^t \setminus \mathcal{I}_i^t \neq \varnothing$, the resulting information structures are of double-sided information asymmetry. Both parties acquire private information hidden from the other, and neither achieves information superiority.

Definition 5 (Incomplete and Imperfect Information).

For the player i, the AIMG is of incomplete information if $\omega \notin \mathcal{I}_i^t$ for all t. The AIMG is of imperfect information if there exists a t such that $\mathcal{I}_i^t \setminus \{\omega\} \subsetneq \mathcal{H}^t \setminus \{\omega\}$.

The following uses lateral movement in 5G networks as a running example to illustrate these information structures in ZTD, which is based on [47].

5.2 Defending against Lateral Movement: A Running Example

Consider a 5G network represented by a directed graph $G = \langle V, E \rangle$, where V is the set of nodes, each of which represents a device/facilities connected to the network, and $E = \{(u, v) | u, v \in V\}$ denotes the set of edges, with each directed edge representing the stored service connection. For example, (u, v) indicates that the user visiting node u can move towards node v using stored credentials. In this example, we assume that the attacker moves laterally using stolen credentials in the 5G network, attempting to reach a sensitive target node with access to some entry node such as mobile devices. The defender aims to validate the user's authentication when accessing neighboring nodes and reject the malicious attacker. This validation can be achieved by Multi-factor Authentication (MFA) [48]. However, Each MFA over the edge incurs a cost, as MFA consumes additional security resources and time that degrade the system performance of the underlying network. The defense objective is to balance the system performance and security by strategically picking a set of edges for authentication validation.

To demonstrate the expressive power of AIMG, we formulate the above defense problem using game-theoretic language developed in Definition 1. Two decision-makers are involved in this game: the defender and the user of an uncertain type. The user's type space is binary $\Omega = \{0, 1\}$, where $\omega = 0$ indicates that the user is legitimate, whereas the user is the malicious attacker if $\omega = 1$. The type distribution ρ can be considered uniform since the two types are indistinguishable from the defender's viewpoint at the beginning. With historical data, the defender can treat the empirical frequency of malicious users as the type distribution, which reflects the defender's prior knowledge of the adversarial environment.

Suppose the attacker visits a node u at time t. Let V^t be the set of neighboring nodes that can be reached using stored credentials. Mathematically, for any $v \in V^t$, there exists a $(u, v) \in E$. Denote the collection of such edges by E^t, and the resulting

subgraph $G^t = \langle V^t, E^t \rangle \subset G$ is referred to as the authentication graph. The user can easily visit any node within the authentication graph if the defender does not impose MFA on E^t. Define $L^t : V^t \rightarrow \{0, 1\}$ as the indicator function. For any $v \in V^t$, $L^t(v) = 1$ is v has been visited before time t, otherwise $L^t(v) = 0$. With a slight abuse of notation, we treat $L^t \in \{0, 1\}^{|V^t|}$ as a binary vector of time-varying dimensions.

The state variable comprises the authentication graph and the indicator, $s^t = (G^t, L^t)$, which captures the progress of the lateral movement and is fully observable to the attacker and the defender. With modern security machinery such as Intrusion Detection System (IDS) [43] and Security Information and Event Management (SIEM) [44], the trace of the user/attacker creates a sequence of events that can be used for security analysis. Consequently, the defender can acquire additional observation of the network system, which is captured by the partial observation o_D in AIMG. The security machinery producing such observation corresponds to the observation function σ_D in Definition 1. Note that the attacker's partial observation is degenerate in this case, i.e., $O_A = \varnothing$.

The action sets of the two parties are specified below. The attacker moves laterally in the network and chooses the next node to visit at each time step. Given the current state $s^t = \langle G^t, L^t \rangle$, the attacker's action set includes a collection of edges $\mathcal{A}_A := \{(u, v) | (u, v) \in E^t, L^t(u) = 1, L^t(v) = 0\}$, of which the outbound node v is to be visited. In APT, the stealthy attacker only picks one edge at each time step to evade detection. To combat the lateral movement, the defender strategically picks a subset of E^t and imposes MFA validation accordingly. Mathematically, the defense action set amounts to the power set of E^t, i.e., the set of all possible subsets of E^t, which is denoted by $\mathcal{A}_D = 2^{E^t}$.

The system evolution is determined by the joint action of both parties, where the attacker picks an edge a_A^t while the defender selects a subset of edges for MFA a_D^t. Given the current authentication graph G^t, one needs to satisfy the MFA requirements if $a_A^t = a_D^t$ before moving to the next node. It is assumed that the legitimate user ($\omega = 0$) has a higher chance to pass this MFA, while the malicious attacker is rejected. On the occasion that $a_A^t \notin a_D^t$, both types can easily move forward. The authentication graph and the visiting history shall be updated accordingly when the user/attacker reaches a new node, and this procedure repeats until the end of the horizon. The horizon length $H \in (0, \infty)$ denotes the maximum time for the attacker to operate within the network without credential renewal. The identity life-cycle lasts for H time steps, after which the stored credentials expire, and the attacker loses the foothold in the network.

The utility function captures the trade-off between operation costs resulting from authentication and system security. From the defender's stance, the cost of authentication validation over an edge is given by the scalar $c : E \rightarrow \mathbb{R}$, and the total cost of imposing MFA on a subset of edges a_D is defined as (with abuse of notation) $c(a_D) = \sum_{e \in a_D} c(e)$. In addition to the authentication cost, system security is also a key factor in the evaluation of defense effectiveness. Denote by v^* the target node, and the indicator function $L^t(v^*)$ implies whether the target has been reached or not. Only when the malicious attacker ($\omega = 1$) visits v^*, the network system is compro-

mised, incurring a devastating cost M. Consequently, the defender's utility depends on the hidden type and is defined below.

$$u_D(s^t, a_D^t, a_A^t, \omega) = \begin{cases} c(a_D^t), & \text{if } \omega = 0, \\ c(a_D^t) + ML^t(v^*), & \text{otherwise.} \end{cases}$$

Likewise, the attacker's utility function is also type-dependent. For the malicious attacker, passing the MFA is laborious and incurs a huge cost $-\hat{M}$. In contrast, the MFA validation is effortless. Whatever the type is, the attacker/user is rewarded by R when arriving at the target node, and they share the same transition cost $u(s^t, a_A^t)$ when navigating within the network. Using mathematical terms, the utility function is as below.

$$u_A(s^t, a_D^t, a_A^t, \omega) = \begin{cases} u(s^t, a_A^t) - RL^t(v^*), & \text{if } \omega = 0, \\ u(s^t, a_A^t) + \hat{M}\mathbb{1}_{\{a_A^t \in a_D^t\}} - RL^t(v^*), & \text{otherwise.} \end{cases}$$

5.3 Trust Evaluation and Access Policy in Zero-Trust Defense

Heretofore, our discussions have primarily concerned the theoretical underpinning of ZTD provided by the game-theoretic framework (AIMG) and AIMG's expressivity regarding information structures. This subsection shifts the focus from ZTD modeling to ZTD design, and the key message is that the game-theoretic solution concept leads to effective and automated ZTD in 5G networks.

We begin with the trust engine and trust evaluation in ZTD. Depending on its architecture, the trust engine can be categorized into attribute-based, Bayesian, and machine-learning-based trust engines. The attribute-based trust engine (ABTE) evaluates the trustworthiness of entities based on their specific attributes or characteristics. Attributes are specific properties or qualities of an entity that are relevant to determining trust, which can include factors such as the security posture of devices and endpoints, the user's location, time of access, and the sensitivity of the requested resource. The evaluation process involves assigning weights or importance to different attributes based on their significance in determining trust. These weights or importance are often pre-defined policies or algorithms, and hence, ABTE relies heavily on the domain knowledge of the security context and involves handcrafting.

The following subsections introduce another two trust engine architectures built upon Bayesian inference and machine learning, leading to automated dynamic trust evaluation capable of adapting to a variety of security scenarios. We refer to the two trust engines as the Bayesian trust engine (BTE) and the machine-learning-based trust engine (MLTE), respectively. A summary of these trust engines is presented in Table 1.

Table 1: A comparison of three kinds of trust engines. Compared with ABTE, BTE and MLTE can adapt to new scenarios without significantly resetting the engine configuration. MLTE is a data-driven trust engine that does not require a complete grasp of the domain knowledge, yet, the price to pay is that its offline pre-training needs a decent amount of data

	Domain Knowledge	Offline Training	Online Computation	Adaptation
ABTE	✓	✗	✗	✗
BTE	✓	✗	✓	✓
MLTE	✗	✓	✓/✗	✓

5.3.1 Bayes Trust Engine

Definition 6 (Bayes Trust Engine). A trust engine is said to be Bayesian if the trust evaluation is produced recursively using the Bayes rule. Let $\mathcal{I}_i^{t+1} \setminus \mathcal{I}_i^t$ be the emerging information at time $t + 1$, then the trust b^{t+1} is obtained by (3a) and the Bayesian update is given by (3b), where $\mathbb{P}(\mathcal{I}_i^{t+1} \setminus \mathcal{I}_i^t | \omega)$ is the probability of observing $\mathcal{I}_i^{t+1} \setminus \mathcal{I}_i^t$ conditional on the hidden type ω.

$$b^{t+1} = \Phi(\mathcal{I}_i^{t+1}) = \Phi(\mathcal{I}_i^{t+1} \setminus \mathcal{I}_i^t, b^t), \tag{3a}$$

$$b^{t+1}(\omega) = \frac{b^t(\omega)\mathbb{P}(\mathcal{I}_i^{t+1} \setminus \mathcal{I}_i^t | \omega)}{\sum_{\omega' \in \Omega} b^t(\omega')\mathbb{P}(\mathcal{I}_i^{t+1} \setminus \mathcal{I}_i^t | \omega')}. \tag{3b}$$

Using the lateral movement example in Section 5.2, the emerging information for the defender at time $t + 1$ is $\mathcal{I}_i^{t+1} \setminus \mathcal{I}_i^t = \{a_A^t, a_D^t, s^{t+1}, o^{t+1}\}$. Given the two parties' policies π_A and π_D, the conditional probability is defined as $\mathbb{P}(a_A^t, a_D^t, o^t, s^{t+1} | \omega) = P(s^{t+1} | s^t, a_D^t, a_A^t, \omega)\sigma(o^t | s^t, a_A^t, a_D^t, \omega)\pi_D(a_D^t | s^t)\pi_A(a_A^t | s^t, \omega)$. Consequently, the belief update is obtained through the following equation.

$$b^{t+1}(\omega) = \frac{b^t(\omega)P(s^{t+1} | s^t, a_D^t, a_A^t, \omega)\sigma(o^t | s^t, a_A^t, a_D^t, \omega)\pi_A(a_A^t | s^t, \omega)}{\sum_{\omega'} b^t(\omega')P(s^{t+1} | s^t, a_D^t, a_A^t, \omega')\sigma(o^t | s^t, a_A^t, a_D^t, \omega')\pi_A(a_A^t | s^t, \omega')}. \tag{4}$$

Compared with the ATE, the BTE adapts to the online environment by processing emerging information recursively without pre-training or preparation. As a plug-and-play engine, BTE requires a decent understanding of the network operation to compute the conditional probability $\mathbb{P}(\mathcal{I}_i^{t+1} \setminus \mathcal{I}_i^t)$, including the system transition P, the security monitoring machinery σ, and the attacker's strategy π_A.

Several remarks are in order on the practicability of BTE. Except for the anticipated strategy π_A, the system transition function P and the observation function σ are readily accessible to the defender. In the lateral movement example, the sys-

tem transition is deterministic: if one edge (u, v) is picked, the next node must be the head node v, and the associated authentication graph and the indicator are determined accordingly. Consider the IDS as the observation function. The corresponding observation space is binary $O_D = \{0, 1\}$, where 0 means no alarm is raised while 1 indicates that a security alert is signaled, warning the defender that the user is more likely to be malicious. In this case, $\sigma(o^t = 1|s^t, a^t, a_D^t, \omega = 1)$ is the detection rate, and $\sigma(o^t = 1|s^t, a^t, a_D^t, \omega = 0)$ is the false alarm rate, both of which are included in the IDS configuration revealed to the defender. As one can see from (4), the attacker's strategy π_A is involved in the Bayesian update, even though it is explicitly included in the information structure $\mathcal{I}_D^t$. Due to the predictive nature of equilibrium in game theory, the defender is able to derive the attacker's optimal strategy using the game tuple in Definition 1, from which the attacker has no incentive to deviate. Using plain words, the defender can anticipate the attacker's strategy π_A and use this predicted strategy the update the trust. Section 5.3.3 elaborates on this equilibrium notion in detail, where we articulate the close connection between BTE and Bayesian Nash equilibrium in game theory, leading to an adaptive zero-trust defense in contrast to ATE.

One computational hurdle of BTE lies in that the denominator in (3b) is given by an integration (summation) of the conditional probability $\mathbb{P}(\mathcal{I}_i^{t+1} \setminus \mathcal{I}_i^t | \omega')$ with respect to the trust $b^t(\omega')$. As the arms race between the defender and the attacker heats up, the attack techniques develop day and night, and consequently, the number of attack types grows astronomical. As a result, the trust evaluation process in the online execution is burdened with great computation overhead, causing authentication latency in ZTD.

In addition to the computation overhead, another limitation of BTE is that it relies heavily on the domain knowledge of the underlying network. Take the lateral movement defense as an example. The observation function σ corresponds to a network security machinery (e.g., SIEM) that monitors the attacker's activities and reports incidents to network operators. Note that such feedback from the security machinery may not be directly applicable in BTE on some occasions since mathematically σ needs to be a conditional probability measure in BTE as shown in (4). For example, if the observation variable $o \in O$ is a log message or an audit trail of the network system, then one needs to infer the attack type distribution behind these security events, requiring certain expertise in network security.

5.3.2 Machine Learning Trust Engine

To address these limitations of BTE, one alternative approach is to utilize machine learning methodologies, which offer an **end-to-end** trust evaluation. The machine-learning-based trust engine undergoes an offline training process before the online execution, and no heavy computation is involved in the online phase, although lightweight model updates can happen on some occasions to adapt the machine-learning model to new security scenarios [47]. Powered by recent advancements in large language models [49] and other related deep learning architectures [50, 51], ML models capable of processing multi-modal inputs (texts and audio, etc.) display

great potential in creating end-to-end trust evaluation that maps the raw system log files to a trust metric without much human involvement. Compared with BTE, MLTE does not require domain knowledge or online computation, yet the price to pay is the pre-training process, and collecting high-quality training data can be cumbersome. This is because the training data shall include incidence reports, system logs, and other related log messages, which often contain sensitive information regarding the network systems, and hence they are not open-sourced. Even if they are, these data come from a specific scenario, and the resulting trust engine may not generalize well to other network defense problems.

Despite its limitations, MLTE provides a data-driven trust evaluation that is suitable for large-scale complex 5G networks. Mathematically, MLTE performs a statistical inference task where the engine infers the hidden type using the observations. The following takes variational Bayes inference (VB) as an example to illustrate how to train and deploy an inference network as the trust engine. In statistical inference, VB refers to a family of techniques in Bayesian inference for approximating the posterior probability of unobserved variables (e.g., hidden types) conditional on the observed ones (e.g., those in the $\mathcal{I}_i^t$). We pick VB because of its close connection with BTE and wide applications in machine learning problems, such as variational autoencoders, which gives rise to many off-the-shelf ML toolsets readily available to network security practitioners. We refer the reader to [51] for more details on statistical inference and its applications.

For simplicity, we drop the time index in the information structure and use $\mathcal{I}$ in the following discussion. Adopting a probabilistic viewpoint, we consider $\mathcal{I}$ and ω as two random variables generated by some random process. The process consists of two steps: 1) a realization ω is generated from the prior ρ; 2) a realization $\mathcal{I}$ is generated from a conditional distribution $\mathbb{P}(\mathcal{I}|\omega)$, which is in a similar vein as (3b). The goal of the inference task is to derive the posterior distribution $\mathbb{P}(\omega|\mathcal{I})$ characterized by the Bayesian rule: $\mathbb{P}(\omega|\mathcal{I}) = \mathbb{P}(\mathcal{I}|\omega)\rho(\omega)/\int \mathbb{P}(\mathcal{I}|\omega)\rho(\omega)d\omega$. Similar to the computation issue in BTE, the integral is intractable.

Denote by $q_\phi(\omega|\mathcal{I})$ a neural network (with parameter $\phi \in \mathbb{R}^n$) approximation to the true posterior $\mathbb{P}(\omega|\mathcal{I})$. Taking inspiration from the evidence lower bound (ELBO) method [51], we derive a loss function for the training purpose whose minimizer $q_{\phi^*}(\omega|\mathcal{I})$ serves as the trust engine in ZTD. Given a realization $\mathcal{I}$, its marginal likelihood can be written as

$$\log \mathbb{P}(\mathcal{I}) = D_{KL}[q_\phi(\omega|\mathcal{I})||\mathbb{P}(\omega|\mathcal{I})] + \mathcal{L}(\phi; \mathcal{I}), \tag{5}$$

where $\mathcal{L}(\phi; \mathcal{I}) = \log \mathbb{P}(\mathcal{I}) - D_{KL}[q_\phi(\omega|\mathcal{I})||\mathbb{P}(\omega|\mathcal{I})]$. $D_{KL}[q_\phi(\omega|\mathcal{I})||\mathbb{P}(\omega|\mathcal{I})] := \mathbb{E}_{q_\phi(\omega|\mathcal{I})}[\log(q_\phi(\omega|\mathcal{I})/\mathbb{P}(\omega|\mathcal{I}))]$ is the KL divergence between the two distributions. The intuition behind this likelihood expression is that the KL divergence $D_{KL}[q_\phi(\omega|\mathcal{I})||\mathbb{P}(\omega|\mathcal{I})]$ in (5) measures the discrepancy between the true posterior $\mathbb{P}(\omega|\mathcal{I})$ and its neural network approximation $q_\phi(\omega|\mathcal{I})$, which is to be minimized. From (5), minimizing the KL term is equivalent to maximizing $\mathcal{L}(\phi; \mathcal{I})$. Since the KL term is non-negative, $\mathcal{L}(\phi; \mathcal{I})$ lower bounds the log-likelihood on the left-hand side, which is referred to as the evidence (or variational) lower bound.

Compared with the KL term $D_{KL}[q_\phi(\omega|\mathcal{I})||\mathbb{P}(\omega|\mathcal{I})]$, this lower bound, rewritten as below, does not explicitly involve the posterior distribution $\mathbb{P}(\omega|\mathcal{I})$. The rest of this subsection is devoted to the stochastic optimization problem $\max_\phi \mathcal{L}(\phi; \mathcal{I})$, which amounts to the pre-training of MLTE.

$$
\begin{aligned}
\mathcal{L}(\phi; \mathcal{I}) &= \log \mathbb{P}(\mathcal{I}) - D_{KL}[q_\phi(\omega|\mathcal{I})||\mathbb{P}(\omega|\mathcal{I})] \\
&= \log \mathbb{P}(\mathcal{I}) - \mathbb{E}_{q_\phi(\omega|\mathcal{I})}[\log q_\phi(\omega|\mathcal{I}) - \log \mathbb{P}(\omega|\mathcal{I})] \\
&= \mathbb{E}_{q_\phi(\omega|\mathcal{I})}[\log \mathbb{P}(\mathcal{I})] - \mathbb{E}_{q_\phi(\omega|\mathcal{I})}[\log q_\phi(\omega|\mathcal{I}) - \log \mathbb{P}(\omega|\mathcal{I})] \\
&= \mathbb{E}_{q_\phi(\omega|\mathcal{I})}[-\log q_\phi(\omega|\mathcal{I}) + \log \mathbb{P}(\mathcal{I}, \omega)].
\end{aligned}
\tag{6}
$$

Consider some dataset $\mathcal{D} := \{\mathcal{I}^{(k)}\}_{k=1}^K$ consisting of K independently identically distributed (i.i.d.) sample observations under random attack types $\omega^{(k)}$ drew from $\rho(\cdot)$. $\mathcal{I}^{(k)}$ represents historical security incidence reports during the network operation, and the superscript (k) denotes the sample index rather than the time step. Note that only the dataset $\mathcal{D}$ is available in training, whereas the variable $\omega^{(k)}$ remains hidden (the prior ρ is known), as often witnessed in real-world scenarios.

In addition to the inference network $q_\phi(\omega|\mathcal{I})$, we introduce a generative network $p_\theta(\mathcal{I}|\omega)$, $\theta \in \mathbb{R}^m$, which approximates the conditional probability $\mathbb{P}(\mathcal{I}|\omega)$. Consequently, the joint distribution $\mathbb{P}(\mathcal{I}, \omega)$ in (6) can also be parameterized: $\mathbb{P}(\mathcal{I}, \omega) = \rho(\omega)p_\theta(\mathcal{I}|\omega)$. With a slight abuse of notation, we denote such parameterization by $p_\theta(\mathcal{I}, \omega)$. Similar to our argument in justifying the use of π_A in (4), $p_\theta(\mathcal{I}|\omega)$ can be interpreted as the defender's conjecture of the attack strategy that eventually leads to the resulting observation $\mathcal{I}$. With this additional parameterization, the lower bound under the datapoint $\mathcal{I}^{(k)}$ becomes

$$
\mathcal{L}(\phi, \theta; \mathcal{I}^{(k)}) = \mathbb{E}_{q_\phi(\omega|\mathcal{I}^{(k)})}[-\log q_\phi(\omega|\mathcal{I}^{(k)}) + \log p_\theta(\mathcal{I}^{(k)}, \omega)].
\tag{7}
$$

The remaining task is simply to approximate the gradient of the expectation in (7) using samples and to apply stochastic gradient descent. Note that the expectation is taken with respect to the hidden variable ω conditional on $\mathcal{I}^k$. Hence, one needs to first draw a batch of M samples $\{\omega^{(k,l)}\}_{l=1}^M$ from q_ϕ, and then compute the gradient estimators

$$
\begin{aligned}
\widehat{\nabla}_\phi \mathcal{L}(\phi, \theta; \mathcal{I}^{(k)}) = &-\frac{1}{M} \sum_{l=1}^K \log q_\phi(\omega^{(k,l)}|\mathcal{I}^{(k)}) \nabla_\phi \log q_\phi(\omega^{(k,l)}|\mathcal{I}^{(k)}) \\
&+ \frac{1}{M} \sum_{l=1}^K \log p_\theta(\mathcal{I}^{(k)}, \omega^{(k,l)}) \nabla_\phi \log q_\phi(\omega^{(k,l)}|\mathcal{I}^{(k)}).
\end{aligned}
\tag{8a}
$$

$$
\widehat{\nabla}_\theta \mathcal{L}(\phi, \theta; \mathcal{I}^{(k)}) = \frac{1}{M} \sum_{l=1}^K \nabla_\theta \log p_\theta(\mathcal{I}^{(k)}, \omega^{(k,l)}).
\tag{8b}
$$

The first gradient estimation in (8a) rests on a Monte Carlo (MC) estimation trick detailed below. The key message of this trick is that the gradient of an expectation

can be expressed as an expectation of another gradient, which can be approximated using Monte Carlo sampling. Suppose, for the time being, one needs to estimate the gradient $\nabla_\phi \mathbb{E}_{q_\phi(\omega)}[f(\omega)]$ where $\mathcal{I}$ is suppressed, and $f(\omega)$ is an arbitrary function. Rewriting the gradient term in the integral form, we obtain

$$
\begin{aligned}
\nabla_\phi \mathbb{E}_{q_\phi(\omega)}[f(\omega)] &= \nabla_\phi \int f(\omega) q_\phi(\omega) d\omega \\
&= \int f(\omega) \nabla_\phi q_\phi(\omega) d\omega \\
&= \int f(\omega) \frac{\nabla_\phi q_\phi(\omega)}{q_\phi(\omega)} q_\phi(\omega) d\omega \\
&= \int f(\omega) \nabla_\phi \log q_\phi(\omega) q_\phi(\omega) d\omega \\
&= \mathbb{E}_{q_\phi(\omega)}[f(\omega) \nabla_\phi \log q_\phi(\omega)].
\end{aligned}
\tag{9}
$$

Therefore, the MC estimation under K samples $\{\omega^{(l)}\}_{l=1}^K$, denoted by $\widehat{\nabla}_\phi$, is given by $\widehat{\nabla}_\phi = 1/K \sum_{l=1}^K f(\omega^{(l)}) \nabla_\phi \log q_\phi(\omega^{(l)})$.

We apply this trick to derive the first gradient estimation. As one can see from the (10), the gradient $\nabla_\phi \mathcal{L}(\phi, \theta; \mathcal{I}^{(k)})$ comprises three terms.

$$
\begin{aligned}
&\nabla_\phi \mathcal{L}(\phi, \theta; \mathcal{I}^{(k)}) \\
&= \nabla_\phi \mathbb{E}_{q_\phi(\omega|\mathcal{I}^{(k)})}[-\log q_\phi(\omega|\mathcal{I}^{(k)}) + \log p_\theta(\mathcal{I}^{(k)}, \omega)] \\
&= \nabla_\phi \int \left(-\log q_\phi(\omega|\mathcal{I}^{(k)}) + \log p_\theta(\mathcal{I}^{(k)}, \omega)\right) q_\phi(\omega|\mathcal{I}^{(k)}) d\omega \\
&= -\underbrace{\int \nabla_\phi \log q_\phi(\omega|\mathcal{I}^{(k)}) q_\phi(\omega|\mathcal{I}^{(k)}) d\omega}_{①} - \underbrace{\int \log q_\phi(\omega|\mathcal{I}^{(k)}) \nabla_\phi q_\phi(\omega|\mathcal{I}^{(k)}) d\omega}_{②} \\
&\quad + \underbrace{\int \log p_\theta(\mathcal{I}^{(k)}, \omega) \nabla_\phi q_\phi(\omega|\mathcal{I}^{(k)}) d\omega}_{③}.
\end{aligned}
\tag{10}
$$

Since $\nabla_\phi \log q_\phi(\omega|\mathcal{I}^{(k)}) = \nabla_\phi q_\phi(\omega|\mathcal{I}^{(k)})/q_\phi(\omega|\mathcal{I}^{(k)})$, $① = \nabla_\phi \int q_\phi(\omega|\mathcal{I}^{(k)}) d\omega = 0$. Applying the trick to the second and third terms, we arrive at the following equations.

$$
\begin{aligned}
② &= \mathbb{E}_{q_\phi(\omega|\mathcal{I}^{(k)})}[\log q_\phi(\omega|\mathcal{I}^{(k)}) \nabla_\phi \log q_\phi(\omega|\mathcal{I}^{(k)})], \\
③ &= \mathbb{E}_{q_\phi(\omega|\mathcal{I}^{(k)})}[\log p_\theta(\mathcal{I}^{(k)}, \omega) \nabla_\phi \log q_\phi(\omega|\mathcal{I}^{(k)})].
\end{aligned}
$$

Replacing all the expectations in ①, ②, and ③, one obtains the MC estimation in (8a). It should be noted that such MC estimation, though intuitive and straightforward, suffers from high variance [52]. One effective remedy is the reparameterization technique [51], and the key idea is that one can express the random variable as

$\omega = g_\phi(\varepsilon, \mathcal{I})$ (reparameterization), where ε is an auxiliary variable with independent marginal $p(\varepsilon)$. When generating $\omega^{(k,l)}$, one follows the procedure: $\varepsilon^{(l)} \sim p(\varepsilon)$ and $\omega^{(k,l)} = g_\phi(\varepsilon^{(l)}, \mathcal{I}^{(k)})$. For example, when $\omega \sim \mathcal{N}(\mu, \Sigma^2)$ (univariate Gaussian with mean μ and variance Σ), a simple reparameterization is $\omega = \mu + \Sigma\varepsilon$, $\varepsilon \sim \mathcal{N}(0, 1)$. Since this parameterization is beyond the scope of this chapter, we refer the reader to [51] for more details on the reparameterization in VB.

5.3.3 Optimal Access Policy: Approximation and Learning

With the trust evaluation process discussed above, we are ready to articulate the access policy π_D in ZTD. To simplify our exposition, we take BTE as the underlying trust engine, and our argument also applies to other kinds of trust engines. Recall that the defender's goal is to minimize the objective function $\min_{\pi_D \in \Pi_D} \mathbb{E}\left\{\sum_{t=1}^H \mathbb{E}_{\omega \sim b^t}[u_D(s^t, a_D^t, a_A^t, \omega)]\right\}$. With a slight abuse of notation, let $u_D(s^t, a_D^t, a_A^t, b^t) = \mathbb{E}_{\omega \sim b^t}[u_D(s^t, a_D^t, a_A^t, \omega)]$ be the expected utility under the trust b^t. Before articulating how to solve the optimal policy, we first address the solution concept in AIMG, i.e., what is the optimality criterion in this multi-agent decision-making?

In general, what distinguishes a game problem from a single-agent optimization is that players' optimization problems are entangled. In AIMG, the defender's problem is given by $\min_{\pi_D \in \Pi_D} \mathbb{E}[\sum_{t=1}^H u_D(s^t, a_A^t, a_D^t, b^t)]$, where the attacker's actions a_A^t are involved. To see this more clearly, we expand the expectation expression, and the defender's problem becomes

$$\min_{\pi_D \in \Pi_D} \mathbb{E}_{\pi_D, \pi_A, P, \sigma}\left[\sum_{t=1}^H u_D(s^t, a_A^t, a_D^t, b^t)\right]. \tag{11}$$

Hence, when the defender determines the access policy, it must take the attacker's move into account and vice versa. From our early argument in BTE, one can view the defender's optimal policy as the minimizer to (11) under the anticipated attacker's strategy π_A^*, i.e.,

$$\pi_D^* \in \arg\min \mathbb{E}_{\pi_D, \pi_A^*, P, \sigma}\left[\sum_{t=1}^H u_D(s^t, a_A^t, a_D^t, b^t)\right]. \tag{12}$$

Then, the remaining question is how to derive such anticipation. From Nash's seminal work [53], one guiding principle is the unilateral deviation principle, which states that π_A^* is a rational anticipation of the attacker's move if the player has no incentive to unilaterally deviate from such strategy, i.e., π_A^* solves the minimization problem in (13). The pair (π_D^*, π_A^*), given by (12) and (13), constitutes a Nash equilibrium of the AIMG. A formal definition is presented in Definition 7.

$$\pi_A^* \in \arg\min \mathbb{E}_{\pi_D^*, \pi_A, P, \sigma} \left[\sum_{t=1}^{H} u_A(s^t, a_A^t, a_D^t, \omega) \right]. \qquad (13)$$

Definition 7 (Perfect Bayesian Nash Equilibrium). Consider the information-asymmetric game with the objectives of the attacker and the defender defined by (11), (12), and (13). A triple of $\langle \pi_D^*, \pi_A^*, \{b^t\}_{t=1}^{H} \rangle$ is said to be the perfect Bayesian Nash equilibrium of this game if it satisfies

$$\pi_D^*(\cdot|s^t, b^t) \in \arg\min \mathbb{E}_{\pi_D, \pi_A^*, P, \sigma} \left[\sum_{\tau=t}^{H} u_D(s^\tau, a_A^\tau, a_D^\tau, b^\tau) \right], \text{ for any } t \in [H],$$

$$(P1)$$

$$\pi_A^*(\cdot|s^t) \in \arg\min \mathbb{E}_{\pi_D^*, \pi_A, P, \sigma} \left[\sum_{\tau=t}^{H} u_A(s^\tau, a_A^\tau, a_D^\tau, \omega) \right], \text{ for any } t \in [H],$$

$$(P2)$$

$$b^{t+1}(\omega) = \begin{cases} \dfrac{b^t(\omega)\mathbb{P}(\mathcal{I}_D^{t+1}\setminus\mathcal{I}_D^t|\omega)}{\sum_{\omega'\in\Omega} b^t(\omega')\mathbb{P}(\mathcal{I}_D^{t+1}\setminus\mathcal{I}_D^t|\omega')} & \text{if } \mathcal{I}_D^{t+1}\text{is realizable,} \\ \text{an arbitrary probability distribution, otherwise.} \end{cases} \quad (C1)$$

$\mathcal{I}_D^t$ is realizable if there exists ω such that the conditional probability $\mathbb{P}(\mathcal{I}_D^{t+1} \setminus \mathcal{I}_D^t|\omega)$ is strictly greater than zero.

In Definition 7, (P1) and (P2) are refinements of (12) and (13), respectively. When $t = 1$, the refinements coincide with (12) and (13), leading to a Nash equilibrium. What makes the refinements "perfect" is that the $\arg\min$ equations hold for any $t \in [H]$. (P1) and (P2) are referred to as the perfectness conditions in game theory [54], meaning that either player has the incentive to deviate from the equilibrium strategy no matter when (time index t) and where (the state s^t and belief b^t) they start to play AIMG. Finally, the equilibrium in Definition 7 is called Bayesian since the belief is generated in a Bayesian manner. (C1) is referred to as the consistency condition: the belief update shall be compatible with the strategy since π_A^* is involved in the Bayesian update, see (4). In summary, this perfect Bayesian Nash equilibrium (PBNE) is the solution concept considered in the rest of this chapter, and the optimal access policy refers to the equilibrium strategy π_D^* in PBNE.

Solving generic PBNE analytically remains largely an open question, even though recent breakthroughs have shed light on the two-stage Markov game case where the PBNE conditions are rephrased using bilevel-bilinear programming [55]. The rest of this subsection is devoted to the numerical approximation of PBNE. Similar to solving single-agent Markov decision processes where computational methods can be divided into value-based [56, 57] and policy-based [58, 59] approaches, the computation of PBNE (approximately) also follows either value-based, i.e., first approximating the expected utility in (P1) and (P2), or policy-based ones,

i.e., searching for the policy directly. The following presents two representative algorithms from the two categories, respectively.

Belief-Value Iteration We begin with the value-based approach. Recall that the perfectness conditions (P1) and (P2) are an extension of Bellman's principle of optimality [60] to the multi-agent setting. Naturally, one can transplant the value iteration algorithm [60] in dynamic programming to AIMG. However, value iteration operates using backward induction, whereas the belief update is a forward process (Bayesian update). Consequently, one cannot update the value function (i.e., the expected utility) and the belief simultaneously.

A variant of value iteration is proposed in [11] to address the conflict between the value function update and the belief update. The gist is that the updates are performed alternatively: updating the value while fixing the belief and vice versa. We refer to such alternative belief/value updates as belief-value iteration (BVI). Denote by $\mathcal{G}(s, b, u_D, u_A)$ the stage game at the state s under the belief b, where the utility functions are $u_D(s, a_A, a_D, b)$ and $u_A(s, a_A, a_D, \omega)$, $\omega \in \Omega$. Let $\texttt{BayesNash}[\mathcal{G}(s, b, u_D, u_A)]$ be the Bayesian Nash equilibrium operator that takes in the stage game utilities and outputs the equilibrium payoffs $(u_D^*, u_A^*) = \texttt{BayesNash}[\mathcal{G}(s, b, u_D, u_A)]$. The equilibrium payoffs (u_D^*, u_A^*) correspond to the minimum in (P1) and (P2), respectively, with the summations inside the expectations are replaced by the stage game utilities. Mathematically, this equilibrium operator is characterized by bilinear programming [11, 55].

The BVI starts with a belief system initialization $\{b^{(t,0)}\}_{t=1}^{H}$. For the k-th iteration, BVI first fixes the belief system $\{b^{(t,k)}\}_{t=1}^{H}$. The k-th value iteration is given by the backward induction below. For $t = H, H-1, \ldots, 1$,

$$
\begin{aligned}
V_D^{(t,k)}(s, b^{(t,k)}), V_A^{(t,k)}(s) &= \texttt{BayesNash}[\mathcal{G}^{(t,k)}(s, b^{(t,k)})], \\
\mathcal{G}^{(H,k)}(s, b) &= \mathcal{G}(s, b, u_D, u_A), \\
\mathcal{G}^{(t,k)}(s, b) &= \mathcal{G}(s, b, u_D + V_D^{(t+1,k)}, u_A + V_A^{(t+1,k)}),
\end{aligned}
\tag{VI}
$$

where $\mathcal{G}^{(t,k)}$ is referred to as the subgame starting from time t during the k-th iteration, bearing the same spirit of the term "cost-to-go" in MDP [60]. The utility function in this subgame is defined in (14). The attacker's utility $u_A + V_A$ can be defined similarly. We remark that by applying the equilibrium operator $\texttt{BayesNash}$ in (VI), the perfectness conditions in Definition 7 are satisfied, and $V_D^{(H,k)}$ and $V_A^{(H,k)}$ returned by (VI) are the equilibrium payoffs of the two players, respectively, under the belief system $\{b^{t,k}\}_{t=1}^{H}$.

$$
u_D + V_D^{(t+1,k)}(s, a_A, a_D, b^{(t,k)}) = u_D(s, a_A, a_D, b^{(t,k)}) + \mathbb{E}_{s' \sim P}[V_D^{(t+1,k)}(s', b^{(t+1,k)})].
\tag{14}
$$

Given the value functions, the defender's and the attacker's policies can be determined accordingly by solving $\mathcal{G}^{(t,k)}$, and we denote the resulting policies by π_D^k and π_A^k, respectively. To complete the k-th iteration, one needs to update the belief system according to the Bayes rule in (4), which is referred to as belief iteration

(BI) in this context shown in (BI). This belief iteration guarantees the consistency between the policies π_D^k, π_A^k and the belief systems $\{b^{(t,k+1)}\}_{t=1}^H$, as mandated by (C1).

$$b^{(t+1,k+1)}(\omega) = \frac{b^{(t,k)}(\omega)\mathbb{P}_{\pi_D,\pi_A}(s^{t+1}|s^t,\omega)}{\sum_{\omega'} b^{(t,k)}(\omega')\mathbb{P}_{\pi_D,\pi_A}(s^{t+1}|s^t,\omega')}, b^{(1,k+1)}(\omega) = \rho(\omega). \quad \text{(BI)}$$

This interleaved procedure repeats until no significant improvement is observed in the updated value functions. Even though intuitive, BVI does not offer any convergence guarantees since the operator `BayesNash` in general is not a contraction mapping [61]. Even assuming it is, we note that the introduction of (BI) further complicates the analysis, and it remains unclear whether the combination of (VI) and (BI) is a contraction mapping. Yet, it is safe to conclude that shall BVI converge, the resulting policies and the belief system must be a PBNE.

Policy Gradient We now shift the focus from the value-based approach to the policy-based one. For simplicity, we fix the attacker's policy in the sequel and present the policy gradient method [58] in reinforcement learning. The key message is that the defender's optimal policy can be learned from sample trajectories using stochastic gradient descent. Consider the defender's problem in (15) where the attacker's strategy is fixed and suppressed.

$$\min_{\pi_D \in \Pi_D} V_D := \mathbb{E}_{\pi_D,P,\sigma}\left[\sum_{t=1}^H u_D(s^t, a_A^t, a_D^t, b^t)\right]. \tag{15}$$

Suppose the policy is parameterized by a neural network $\pi_D(\phi), \phi \in \mathbb{R}^n$. Then, one can search for the optimal policy through gradient descent, i.e., $\phi \leftarrow \phi - \nabla V_D(\phi)$ (the learning rate is suppressed). $\nabla V_D(\phi) = \nabla\mathbb{E}_{\pi_D(\phi),P,\sigma}[\sum_{t=1}^H u_D(s^t, a_A^t, a_D^t, b^t)]$. Recall the MC estimation trick in (9), we rewrite the gradient as in (16), referred to as the policy gradient.

$$\nabla V_D(\phi) = \mathbb{E}_{\pi_D(\phi),P,\sigma}\left[\nabla \log \pi_D(\phi) \sum_{t=1}^H u_D(s^t, a_A^t, a_D^t, b^t)\right]. \tag{16}$$

Denote a sample trajectory under the policy $\pi_D(\phi)$ (in short, ϕ) by $\ell(\phi) := \{s^1, a_A^1, a_D^1, u_A^1, u_D^1, o^1, \ldots, s^H, a_A^H, a_D^H, u_A^H, u_D^H, o^H\}$, where $u_D^t = u_D(s^t, a_A^t, a_D^t, b^t)$, b^t is derived using the Bayes rule in (4). Then, an unbiased estimate of $\nabla V_D(\phi)$, denoted by $\widehat{\nabla}V(\phi)$ is constructed as $\widehat{\nabla}V(\phi) = \nabla \log \pi_D(\phi) \sum_{t=1}^H u_D^t$. Denote by $u_D(\ell) = \sum_{t=1}^H u_D^t$ the empirical return of the sample trajectory. One common practice to reduce the variance of the MC estimate $\widehat{\nabla}V(\phi)$ is to collect a batch of trajectories $\{\ell^{(k)}\}_{k=1}^K$ and take the average: $\widehat{\nabla}V(\phi) = 1/K \sum_{k=1}^K \nabla \log \pi_D(\phi) u_D(\ell^{(k)})$. Starting from an initialization ϕ^0, one need first implement the policy $\pi_D(\phi^0)$ in a simulated network system [62] and collect a batch of trajectories $\{\ell^{(k)}\}_{k=1}^K$. Then, the policy is updated using the policy gradient discussed above. The procedure repeats until the parameter ϕ^k stabilizes. Since policy gradient is a first-order method, it is only

guaranteed to converge to the first-order stationary point where $\nabla V_D(\phi) = 0$. Even though this first-order point may not be the exact equilibrium point, it often leads to satisfying defense policy, as observed in the literature [63].

5.4 Generalizability, Explainability, and Accountability of Learning-based Zero-Trust Defense

5.4.1 Reinforcement Learning and Explainable Defense

Even though RL leads to a theoretically guaranteed approach to learning the ZTD policy, the missing part is that the learned policy, i.e., the model weights of the neural network, remains a black box and is difficult for human operators to comprehend. The explainability of RL (XRL), as an emerging field devoted to casting light on the inner workings of RL agents, has gained momentum across various research communities. Since XRL is still in its infancy, there is no consensus over the exact definitions of explainability, and most of the current endeavors try to explain the actions of RL agents [64]. Following this line of research, we discuss the explainability of the optimal access policy learned by RL in the following, which addresses the question:

How does the RL policy grant or deny access based on the trust evaluation?

Our XRL approach exploits the mathematical structure of the AIMG and utilizes non-parametric policy learning, i.e., the RL policy is expressed in closed form without involving neural networks [65, 47]. Hence, our XRL study is more aligned with the interpretability of the RL policy, indicating that the intrinsic logic of the defense mechanism is transparent and easy to understand rather than a post-hoc property.

The gist of the explainability in ZTD is that the optimal policy is of a threshold form [65]. Consider the lateral movement case in Section 5.2 as an example, where the type space and the defense action space are binary: $\Omega = \{0, 1\}$ (0-legitimate user, 1-attacker) and $\mathcal{A}_D = \{0, 1\}$ (0-active defense, 1- inactive). In this example, the belief b resides in the two-dimensional probability simplex, which can be uniquely determined by its entry $b(0)$. We refer to $b(0) \in [0, 1]$ as the trust score, implying the likelihood of the user is legitimate. A threshold policy $\pi_D(b)$ is defined in (17), and the threshold is given by τ. As its name suggests, the defense remains idle as long as the trust score is above the threshold, while it is activated once the trust score is below the critical value.

$$\pi_D(b) = \begin{cases} 0, \ 0 \le b(0) \le \tau, \\ 1, \ \tau < b(0) \le 1. \end{cases} \tag{17}$$

The advantage of this threshold policy is self-evident: it is a white box clearly displaying how the trust evaluation is utilized. The same policy gradient method presented above also applies to the learning of thresholds. Even though the gradient $\nabla_\tau \pi_D$ does not acquire a closed form, one can leverage the simultaneous perturbation stochastic approximation (SPSA) to estimate the gradient [65, 47]. The threshold

form in (17) also extends to the finite-action case, where $|\mathcal{A}_D| - 1$ threshold values partition the interval $[0, 1]$ into $|\mathcal{A}_D|$ subintervals (the type space is still binary).

5.4.2 Meta-Learning and Generalizable Defense

The limitation of the threshold policies is concerned with generalization ability. The optimal policy (or equivalently, threshold) trained in one network setup cannot deal with another scenario where the system vulnerabilities are different from the training setup. To facilitate our discussion, denote by $\theta \in \Theta$ the network system configuration that can affect the system transition P (or the observation function σ) under this configuration. Using the notations in Definition 1, the defender now faces a family of games, and the transition function P_θ of each game is parameterized by θ subject to a distribution $p(\theta)$. We refer to each game under parameter θ as an attack scenario. The policy trained for the scenario θ does not generalize well to θ', leading to ineffective ZTD.

To equip ZTD with generalizability under information asymmetry, a scenario-agnostic ZTD (SA-ZTD) is proposed in [47], creating a generalizable ZTD capable of handling new attack scenarios unseen in the training phase. SA-ZTD rests on meta-learning, an emerging learning paradigm that aims to learn a learning strategy using training data [66]. In the face of a new scenario unseen in the training phase, the obtained learning strategy enables the defender to learn a new defense on the fly using far fewer data than from scratch. This idea of defending on the fly is also explored in adversarial machine learning leading to impressive defense performance [67]. Since real-world applications involve a large (possibly infinite) number of attack scenarios, it is intractable to learn the optimal policy for each scenario. Powered by meta-learning, SA-ZTD uses only a handful of known scenarios, more precisely, sample trajectories from these scenarios. Hence, the word "agnostic," whose root means "not known," is used to emphasize that the adaptation ability is acquired without knowledge of the network configuration of every scenario.

Two pillars of SA-ZTD are the meta policy π_{meta} and the adaptation mapping $\Psi : \Pi_D \times \Theta \to \Pi_D$. The adaptation mapping corresponds to the learning strategy mentioned earlier that adapts the meta policy to a new defense $\Psi(\pi_{meta}, \theta)$ when facing a new scenario θ. A formal definition of SA-ZTD is given in [47], which we restate in Definition 8.

Definition 8 (SA-ZTD). A pair $\langle \pi_{meta}, \Psi \rangle$ is said to be a scenario-agnostic zero-trust defense (SA-ZTD) with respect to a scenario distribution $p \in \Delta(\Theta)$ if the pair solves for the minimization problem

$$\min_{\pi, \Psi} \mathbb{E}_{\theta \sim p}[V_D(\Psi(\pi, \theta))]. \tag{18}$$

Similar to empirical risk minimization (ERM) [68, 69], a solution to (18) is obtained by solving the sample average approximation:

$$(\pi_{meta}, \Psi) \in \arg\min \frac{1}{|\widehat{\Theta}|} \sum_{\theta \in \widehat{\Theta}} V_D(\Psi(\pi, \theta)), \tag{19}$$

where $\widehat{\Theta} \subset \Theta$ is a finite collection of scenarios i.i.d. sampled from $p \in \Delta(\Theta)$. The term "agnostic" points to the fact that the exact scenario distribution p is usually unknown in security practice and often replaced by an empirical distribution provided by security datasets, such as the data from MITRE ATT&CK [70] considered in [47]. In summary, the training of SA-ZTD does not explicitly require the domain knowledge of each attack scenario, such as the system configuration and the observation functions.

Since the function class $\{\Psi | \Psi : \Pi_D \times \Theta \to \Pi_D\}$ is infinite-dimensional, directly seeking an adaptation mapping through (18) [or (19)] is intractable. One remedy is to restrict the focus to the parameterization class where the mapping is parameterized by $\gamma \in \mathbb{R}^n$, $n \in \mathbb{Z}_+$. For example, Ψ_γ can be parameterized by recurrent neural networks, where γ is the model weights and the optimal adaptation is determined by training algorithms [71]. Another well-accepted parameterization is the gradient-based adaptation: $\Psi_\gamma(\pi, \theta) := \pi - \gamma \nabla V_D(\pi)$, and γ is the gradient step size to be optimized [72].

To arrive at an explainable SA-ZTD, one can pick the gradient-based adaptation, as it naturally applies to the non-parametric threshold policies discussed in Section 5.4.1. To be consistent with previous notations, we replace π with τ whenever speaking of threshold policies, where the τ denotes the threshold value. The minimization problem in (18) turns into

$$\min_{\tau \in [0,1]} \mathbb{E}_{\theta \sim p} [V_D(\text{Proj}_{[0,1]}\{\tau - \gamma \nabla V_D\})]. \tag{20}$$

The resulting meta policy, as the minimizer to (20), takes the threshold form that is explainable to human operators, increasing the accessibility and transparency of learning-based ZTD. As argued in [47], the policy gradient method is still applicable to (20). Even though the computation expenditure in SA-ZTD is higher than the vanilla RL policy in (15), the meta policy can adapt to a variety of new scenarios without training from scratch.

5.4.3 Accountability

The accountability of machine-learning-based ZTD (ML-ZTD) refers to the responsibility and answerability of those involved in the design, development, deployment, and use of machine learning or artificial intelligence technologies in general. Accountability aims to ensure that ML-ZTD is developed and utilized in a manner that is ethical, transparent, and fair. What distinguishes accountability of ZTD in 5G

networks from other AI systems is the focus on accountability in system engineering, which encompasses three key aspects: responsibility, detectability, and attribution.

Responsibility Accountability rests on the acknowledgment that individuals and organizations involved in ML-ZTD development and deployment have responsibility for the ZTD's behavior and impact on the network system. Specifically, this responsibility revolves around the question of whether each component involved in ZTD architecture, such as the security machinery, the trust engine, and the access policy, contributes to an ethical, transparent, and fair operation in the network. To be more precise, this responsibility provides compliance requirements and failure standards for each component.

Detectability Responsibility gives the rule book, and the next question to address is whether ZTD operation violates the compliance requirements. Mathematically, the detectability question pertains to statistical inference, such as hypothesis testing and VB methods, where one infers the ground truth (violation) from collected data. Yet, ZTD in 5G networks is a game problem, see Definition 1, where the strategic decision-maker can evade the detection, which must be taken into account when inspecting the ZTD operation. Game theory naturally provides a system-science viewpoint on the detectability question in multi-agent systems, where the incentives, capabilities, and private information of the investigator and the investigatee can be captured through the AIMG in Definition 1. This game-theoretic viewpoint leads to a strategic detection framework.

Attribution No node is an island in large-scale complex 5G networks, and one failing node or component may spur a chain reaction over the network and the ZTD system. When facing a cascading failure in the network defense, one needs to identify the root cause and upgrade the ZTD accordingly. One shall not confuse detection with attribution, even though both of them aim to identify the malfunctioning part of the ZTD and the network system. However, detection addresses the question "where it is", whereas attribution focuses on "why it is such." Mathematically, attribution amounts to a causal inference task [59], where the casual relationship among random variables is established using data.

6 Decision-dominance Defense

While ZTD provides us with a comprehensive framework for trust evaluation and access policy, the networked entities still face multi-stage persistent cyber threats. Therefore, it is crucial to adopt an integrated defense approach that recognizes the intrinsic value of the cyber defense chain and the fundamental principles of zero trust. Decision dominance defense (D^3), which conceptualizes the interactions of cyber defense/kill chain as a stochastic process, forms the backbone of the holistic defense mechanism, with zero trust defense acting as a critical component at every stage. By treating the cyber defense chain as a dynamic system, we acknowledge

the unpredictable nature of cyber threats and the need for proactive decision-making based on real-time information. By incorporating zero trust principles throughout this process, from initial access controls to ongoing monitoring and incident response, we create a robust and resilient defense model that embraces uncertainty, eliminates blind spots, and ensures continuous protection against the relentless onslaught of cyber threats.

Understanding the intricacies of an attack is crucial for developing effective defense strategies. A traditional Lockheed Martin Kill Chain [73, 74] usually outlines seven distinct stages that malicious actors typically follow. These stages include Reconnaissance, where attackers gather information on potential targets; Weaponization, where they create malicious tools or payloads; Delivery, the method through which the attack is transmitted; Exploit, where vulnerabilities are leveraged to gain access; Installation, the establishment of a foothold within the target system; Command & Control, the creation of communication channels for remote control; and finally, Actions on Objectives, where the attacker achieves their intended goals within the compromised system. Comprehensively analyzing and understanding each stage of the Kill Chain requires the defender to effectively engage with adversaries while minimizing the time it takes for an attack to unfold. A proactive cyber defense chain (e.g., [75, 76]) aims to disrupt and curtail the attacker's progress at each stage of the Kill Chain, reducing their opportunity to inflict significant damage. D^3 integrates real-time threat intelligence, advanced analytics, and rapid response mechanisms, including monitoring, detection, response, and attribution, maximizing the abilities to mitigate and neutralize the threats, actively impeding the attacker's progress and shortening the overall time it takes for an attack to materialize. It empowers the 5G network defender to take a more active role in their defense, enabling them to stay one step ahead of the adversary and significantly enhance their resilience against evolving cyber threats.

The essence of D^3 is the critical timing of cutting off the cyber kill/defense chain. In MWD scenarios, while the general concept of understanding, deciding, acting, and assessing fast still holds (i.e., strangling the threats in its cradle), one must take the real-time warfare conditions and game-theoretic thinking into consideration, "knowing oneself and knowing the enemy". Therefore, in the sequel, we formalize D^3 as a Dynkin's type of optimal stopping game acting on a Markov chain of multi-stage cyber-attacks/defense [77], and characterize the equilibrium strategy between the two competitive parties. While our model is built upon ZTD components, the notations should not be confused with the previous section.

6.1 D^3 as Dynkin's Game

By convention, let $(\Omega, \mathcal{F}, \mathbb{P})$ be the probability space. Denote the time index during a lifecycle of the interactions between the cyber kill/defense chain by $t = 1, \ldots, T$. Let $(X_t)_{0 \leq t \leq T}$ be a Markov process modeling the cyber threats, living in space $(\mathcal{X}, \mathcal{G})$, and are adapted to the filtration $\mathbb{F} = (\mathcal{F}_t)_{0 \leq t \leq T}$ with transition kernel $\mathcal{P}$. The

Markovian state captures the identifiable elements in the system, e.g., it can represent the Structured Threat Information eXpression language (STIX) that facilitates this effort [77]. The collection of STIX-type data requires active interactions between the two parties.

We are given three payoff functions $\phi, \zeta, \psi : \mathcal{X} \to \mathbb{R}$ that capture the cyber risk given system states, where from the defender's perspective, (the attacker's perspective would be the opposite,)

1. ϕ is the early termination payoff, which is activated when the cyber defender actively terminates the persistent monitoring/detection and resets the system credential before the malicious operations, including data exfiltration, denial of service, and delivery of ransomware, etc. are executed;
2. ψ is the late response payoff, which is activated when the cyber defender responds to the data exploitation and command & control actions without summarizing the monitoring/detection phase.
3. ζ is the confrontation payoff, which is activated when both parties have extracted information through lateral movement/monitoring and engaging, etc., and perform attack/defense actions at the same stages.

It is reasonable to assume that $\min(\psi, \phi) \leq \zeta \leq \max(\psi, \phi)$, since the confrontation often happens when attackers and defenders both have neutralized assessments for the system, it sits in between the worst and best payoffs.

Here, for simplicity, we first consider the case where the information is symmetrical between the network operator/defender and the attacker, i.e., both parties have access to the state and utility information. However, this formalism shall not exclude the cases where the information is asymmetric and/or the utility functions are unknown/uncertain to one of the parties.

On top of the lower-level cyber threats/defense operations, we define stopping times $\tau, \sigma : \Omega \to \{0, \ldots, T\}$ to capture the termination decisions for both parties. With the assumption that both the attacker and the defender have access to the system state X_t, τ, σ are $\mathbb{F}$-measurable. Denote the set of $\mathbb{F}$-stopping times by $\mathcal{T} := \{0 \leq \tau \leq T : \{\tau(\omega) \leq k\} \in \mathcal{F}_k \; \forall k \in [T], \forall \omega \in \Omega\}$. Moreover, we expect there to be a $2^{[T]}/\mathcal{G}$-measurable map $\tau : \mathcal{X} \to [T]$, where $[T] = \{0, \ldots, T\}$, such that the defender/attacker will make termination decisions based on the information extracted from X_t, without awareness of each other's stopping decisions.

For stopping times τ, σ, the value/cost function for the defender/attacker is defined as:

$$V^{\tau,\sigma}(x) = \mathbb{E}_x[H(\tau,\sigma)] = \mathbb{E}_x\left[\phi(X_\tau)\mathbb{1}_{\{\tau<\sigma\}} + \psi(X_\sigma)\mathbb{1}_{\{\tau>\sigma\}} + \zeta(X_\tau)\mathbb{1}_{\{\tau=\sigma\}}\right], \tag{21}$$

where $H(\tau,\sigma) : \mathcal{T} \times \mathcal{T} \times \Omega \to \mathbb{R}$ is the random payoff of stopping strategies τ and σ, $\mathbb{E}_x$ is the conditional expectation operator with respect to the transition kernel $\mathcal{P}_x$, i.e., there is an operator $\mathcal{T}$ that is a semi-group, such that for any $\mathcal{B}(\mathbb{R})/\mathcal{G}$-measurable function g and $t = 0, \ldots, T$,

$$\mathcal{T}^t g(x) := \mathbb{E}_x[g(X_t)] = \underbrace{\int_X \cdots \int_X}_{t \text{ times}} g(x_t) d\mathcal{P}_{x_{t-1}}(x_t) \ldots d\mathcal{P}_x(x_1).$$

In practice, the convolutional integral is hard to compute directly. Instead, we can leverage sampling methods such as Markov Chain Monte-Carlo (MCMC) to approximate the conditional expectation.

Definition 9 summarizes our game-theoretic formalism.

Definition 9 (Decision Dominance Game). A tuple $(X, \mathcal{P}, \phi, \zeta, \psi, \mathcal{T})$ encapsulates a Decision Dominance Game (DDG) if it satisfies the following:

- there exists a Markov process $(X_t)_{0 \le t \le T}$ that lives in $(X, \mathcal{G})$ with transition kernel $\mathcal{P}$, which can be extracted as cyber threats information;
- ϕ, ζ, and ψ are payoff functions mapping from X_t to $\mathbb{R}$, $\phi, \zeta, \psi \in \mathcal{E}(X)$, which is the set of all bounded $\mathcal{B}(\mathbb{R})/\mathcal{G}$-measurable functions on $(X, \mathcal{G})$. Further, $\min(\phi, \psi) \le \zeta \le \max(\phi, \psi)$ on X;
- at each stage t, both parties pick a stopping strategy from space $\mathcal{T}_t := \{t \le \tau \le T : \{\tau(\omega) \le k\} \in \mathcal{F}_k \ \forall k \in [T], \forall \omega \in \Omega\}$ to decide whether to stop or continue the kill/defense chain.
- at each stage the utility function of the defender is

$$H(\tau_t, \sigma_t) = \phi(X_{\tau_t})\mathbb{1}_{\{\tau_t < \sigma_t\}} + \zeta(X_{\tau_t})\mathbb{1}_{\{\tau_t = \sigma_t\}} + \psi(X_{\sigma_t})\mathbb{1}_{\{\tau_t > \sigma_t\}},$$

while the attacker attains $-H(\tau_t, \sigma_t)$.

Figure 6 gives an example of the DDG outcome. The solution concept of a DDG is given in Definition 10.

Definition 10 (Decision-Dominance Equilibrium (DDE)). A pair of stopping time strategies $(\tau^*, \sigma^*) \in \mathcal{T} \times \mathcal{T}$ is a Decision-Dominance Equilibrium (DDE) if for all initial state $x \in X$, it satisfies the minimax condition:

$$
\begin{aligned}
V^{\tau^*, \sigma^*}(x) &= \text{ess sup}_{\sigma \in \mathcal{T}} \, \text{ess inf}_{\tau \in \mathcal{T}} \, V^{\tau, \sigma}(x) \\
&= \text{ess inf}_{\tau \in \mathcal{T}} \, \text{ess sup}_{\sigma \in \mathcal{T}} \, V^{\tau, \sigma}(x).
\end{aligned}
\tag{22}
$$

The existence of such a value function, however, is a non-trivial question, as we are looking for a pure strategy Nash equilibria in an infinite-dimensional space ($\mathcal{T} \times \mathcal{T}$), Von-Neumann's Minimax theorem does not apply here. However, under certain conditions, we are able to show that a DDG with information symmetry always admits a value function, which is unique up to a state-wise constant translation.
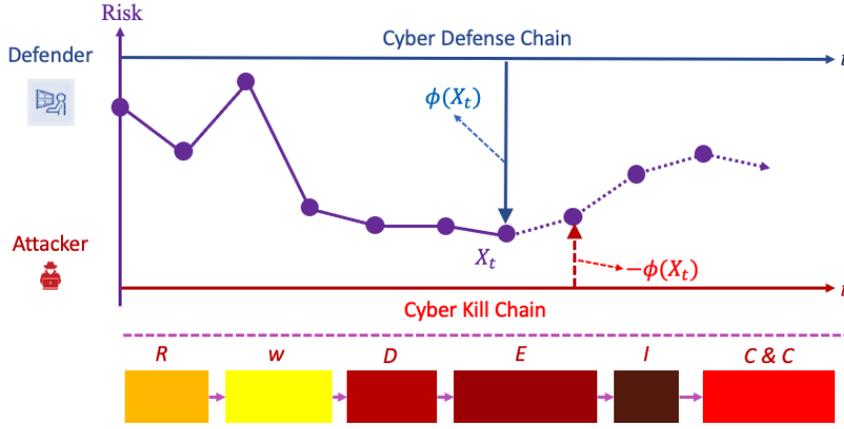
Fig. 6: An illustration of the cyber kill/defense chain interaction. In this case, at time t, the system state has evolved into X_t, the defender cuts off the chain interaction earlier than the attacker and gets payoff $\phi(X_t)$, while the attacker gets $-\phi(X_t)$ since she plans to stop at the next time step

We know from Dynkin's result [78] that when $\phi \leq \zeta \leq \psi$ on $x \in X$, there exists a value process

$$
\begin{aligned}
V_t &= \min\{\psi(X_t), \max\{\phi(X_t), \mathbb{E}[V_{t+1}|\mathcal{F}_t]\}\} \\
&= \max\{\phi(X_t), \min\{\psi(X_t), \mathbb{E}[V_{t+1}|\mathcal{F}_t]\}\},
\end{aligned}
\tag{23}
$$

and the equilibrium strategies capture V_t's hitting times of the upper/lower limits. However, the ordered-payoff assumption is hard to verify in the context of MDW, a more reasonable assumption, as has been discussed before, is $\min(\phi, \psi) \leq \zeta \leq \max(\phi, \psi)$ on X. In addition, the ubiquitous information asymmetry in cyberspace oftentimes makes the derived equilibrium strategies inapplicable.

Therefore, in the sequel, we dive into the more general case defined as in Definition 9, and lay out some essential analytical characterization for the equilibrium value process; further, we give a rough description for the case under information asymmetry.

6.2 Equilibrium Strategies for D^3

In this section, we investigate the existence and characterization of the DDE in two different cases under a symmetric information structure and then discuss an extension. The first case is when the early termination payoff ϕ dominants the late termination payoff ψ, which we call *adversarial dominance*, as in this case, the outcome of engaging in the long term favors the adversary. The second case is called *defense dominance*, where the late termination payoff ψ dominates the early termination payoff ϕ. Hence, the defender is able to endure the kill/defense chain interactions longer than the adversary does.

6.2.1 Case I: Adversarial Dominance

Under the Adversarial Dominance Condition (ADC), the payoff functions satisfy the ordered condition $\psi \leq \zeta \leq \phi$ for all system states $x \in \mathcal{X}$. In this case, at any state $x \in \mathcal{X}$, the defender aims to investigate the kill chain for a proper period of time while trying to terminate the operations faster than the attacker, as it is more costly to wait for the attacker to exploit the vulnerabilities by doing Command & Control than to shut down the service and reset the credentials. This is also called *first-mover advantage*, that is, the defender has the incentive to end the game faster than the opponent.

We shall proceed with the analysis by giving a constructive sequence of equilibrium values. To this end, we investigate the $t \in [T]$ stage problem through backward induction and let $\{V_n^t\}_{n=0}^t$ be the equilibrium processes attained by stopping at no more stage t. At $t \in [T]$, both parties have to choose confrontation, thus at the final stage, the payoff is $\zeta(X_t)$; at $n \in [t-1]$, either they both stop and get payoff value $\zeta(X_n)$, or wait for the next round, in which case the defender has to judge if the termination values $\phi(X_n)$ is higher than the expected engaging values $\mathbb{E}\left[V_{n+1}^t | \mathcal{F}_n\right]$, given that the attacker chooses to engage. Mathematically, we have the value processes for arbitrary $t \in [T]$,

$$
V_t^t = \zeta(X_t),
$$
$$
V_n^t = \text{val} \begin{bmatrix} \zeta(X_n) & \phi(X_n) \\ \psi(X_n) & \mathbb{E}\left[V_{n+1}^t | \mathcal{F}_n\right] \end{bmatrix}, \qquad \text{for } n = t-1, \ldots, 0. \tag{24}
$$

where $\text{val}(\cdot)$ stands for a special value operator of the matrix game, which we interpret as:

$$
V_n^t = \begin{cases} \mathbb{E}\left[V_{n+1}^t | \mathcal{F}_n\right] & \text{if } \phi(X_n) < \mathbb{E}\left[V_{n+1}^t | \mathcal{F}_n\right], \\ \zeta(X_n) & \text{otherwise .} \end{cases}
$$

It turns out that the value processes possess the monotone property (Lemma 1).

Lemma 1. *For every $n, t \in [T]$ such that $n \leq t$, one has that the equilibrium value processes defined as in* (24) *satisfy*

$$V_n^n \leq V_n^t \leq V_n^{t+1}.$$

One can show Lemma 1 with an induction argument. Let the event $E_n := \{\omega : \phi(X_n) < \mathbb{E}[\zeta(X_{n+1})|\mathcal{F}_n]\}$ be when the next round expected confrontational payoff is higher than the current early termination payoff. Consider the base case; it follows that at any stage $t \in [T-1]$, since the next round both parties need to terminate, it is reasonable for the defender to choose to terminate if the early termination payoff is higher than the expected confrontational payoff. Thus,

$$
\begin{aligned}
V_t^{t+1} &= \begin{cases} \mathbb{E}[\zeta(X_{t+1})|\mathcal{F}_t] & \text{on } E_t, \\ \zeta(X_t) & \text{on } E_t^c, \end{cases} \\
&\geq \begin{cases} \phi(X_t) & \text{on } E_t, \\ \zeta(X_t) & \text{on } E_t^c, \end{cases} \\
&\geq \zeta(X_t) = V_t^t.
\end{aligned}
$$

Now we assume that $V_j^{j+k-1} \leq V_j^{j+k}$ for some arbitrary stage $1 \leq k \leq T-1$ and for all $j \in [T-k]$, then, for $t \in [T-k-1]$,

$$
\begin{aligned}
V_t^{t+k} &= \text{val} \begin{bmatrix} \zeta(X_t) & \phi(X_t) \\ \psi(X_t) & \mathbb{E}\left[V_{t+1}^{t+k}|\mathcal{F}_t\right] \end{bmatrix} \\
&\leq \text{val} \begin{bmatrix} \zeta(X_t) & \phi(X_t) \\ \psi(X_t) & \mathbb{E}\left[V_{t+1}^{t+k+1}|\mathcal{F}_t\right] \end{bmatrix} \\
&= V_t^{t+k+1}.
\end{aligned}
$$

Hence, the monotonicity follows by the induction argument.

That V_k^t being increasing in t gives off two signals; the first is that due to the Monotone Convergence theorem for $\mathbb{E}[\cdot|\mathcal{F}_t]$, there exists a limit for V_k^t if we consider the infinite-stage problem ($t \to \infty$); the second is that the dominating strategy can be obtained when the stopping stage is not constrained, up to time T.

Now we define two stopping times, for $t \in [T]$,

$$
\begin{aligned}
\bar{\tau}_t &= \inf\{t \leq k \leq T | V_k^T = \zeta(X_k)\}, \\
\bar{\sigma}_t &= \inf\{t \leq k \leq T | V_k^T = \zeta(X_k)\}.
\end{aligned}
$$

The significance of $(\bar{\tau}_t, \bar{\sigma}_t)$ is given in Theorem 1.

Theorem 1. *Under ADC, the following statements hold for arbitrary initial state $x \in \mathcal{X}$:*

i) For every $t \in [T]$, and all $\tau \in \mathcal{T}_t, \sigma \in \mathcal{T}_t$,

$$\mathbb{E}[H(\tau, \bar{\sigma}_t)|\mathcal{F}_t] \leq V_t^T = \mathbb{E}[V_{\bar{\tau}_t \wedge \bar{\sigma}_t}^T|\mathcal{F}_t] = \mathbb{E}[H(\bar{\tau}_t, \bar{\sigma}_t)|\mathcal{F}_t] \leq \mathbb{E}[H(\bar{\tau}_t, \sigma)|\mathcal{F}_t].$$

ii) At every time $t \in [T]$, a pair $(\bar{\tau}_t, \bar{\sigma}_t)$ is an equilibrium point for that time step t, and a DDE value corresponding to $(\bar{\tau}_0, \bar{\sigma}_0)$ is given as

$$\mathbb{E}[V_0^T] = \mathbb{E}[V_{\bar{\tau}_0 \wedge \bar{\sigma}_0}^T] = \mathbb{E}[H(\bar{\tau}_0, \bar{\sigma}_0)].$$

Proof. Fix a $t \in [T]$ arbitrarily. We have that, if $k \in \{t, \ldots, \bar{\sigma}_t\}$, by definition of $\bar{\sigma}_t$, we have

$$V_k^T = \mathbb{E}[V_{k+1}^T | \mathcal{F}_k].$$

Thus, the sequence $\{V_{k \wedge \bar{\sigma}_t}^T, k \geq t\}$ is a regular Martingale, so that $V_t^T = \mathbb{E}[V_{\tau \wedge \bar{\sigma}_t}^T | \mathcal{F}_t]$ for any $\tau \in \mathcal{T}_t$, by Doob's optional sampling theorem. Since $V_{\bar{\sigma}_t}^T = \zeta(X_{\bar{\sigma}_t}) \geq \psi(X_{\bar{\sigma}_t})$, if $\bar{\sigma}_t \leq \infty$ and $V_k^T \geq \phi(X_k)$ if $\bar{\sigma}_t > k$, it follows that:

$$
\begin{aligned}
V_t^T &= \mathbb{E}[V_{\tau \wedge \bar{\sigma}_t}^T | \mathcal{F}_t] \\
&= \mathbb{E}[V_\tau^T \mathbb{1}_{\{\tau < \bar{\sigma}_t\}} + V_{\bar{\sigma}_t}^T \mathbb{1}_{\{\bar{\sigma}_t \leq \tau\}} | \mathcal{F}_t] \\
&\geq \mathbb{E}[\phi(X_\tau) \mathbb{1}_{\{\tau < \bar{\sigma}_t\}} + \psi(X_{\bar{\sigma}_t}) \mathbb{1}_{\{\bar{\sigma}_t < \tau\}} + \zeta(X_{\bar{\sigma}_t}) \mathbb{1}_{\{\bar{\sigma}_t = \tau\}} | \mathcal{F}_t] \\
&= \mathbb{E}[H(\tau, \bar{\sigma}_t) | \mathcal{F}_t].
\end{aligned}
$$

A symmetric argument can be applied to prove the $\leq$ side for all $\sigma \in \mathcal{T}_t$. By letting $t = 0$ we arrive at the conclusion.

Theorem 1 i) implies that for every subgame starting from time t, the equilibrium strategy is always a threshold strategy for both parties, where the threshold needed to be computed is $\mathbb{E}[V_{t+1}^T | \mathcal{F}_t]$. Both parties have incentives to stop only when $\phi(X_t)$ is hitting the threshold. ii) states that in the adversarial dominance environment, $(\bar{\tau}_0, \bar{\sigma}_0)$ are the equilibrium strategies. However, the determination of the equilibrium value sequence V_t^T is computationally intractable, as one would have to construct the random variables backwardly according to (24), enumerating over the filtration sets.

Therefore, it is crucial to generalize the above arguments to the space of $\mathcal{E}(X)$. As we may assume that the players have access to the payoff functions, constructing a map between X_t and the equilibrium value process can be relatively easier. Indeed, due to the Markovian property of X_t, it turns out we only need a sequence of $\mathcal{B}(\mathbb{R})/\mathcal{G}$-measurable value functions $\{v_t(\cdot)\}_{t \in [T]}$ that satisfies the following conditions:

$$
\begin{aligned}
v_T(x) &= \zeta(x), \quad \text{for all } x \in X, \\
v_t(x) &\in \text{SE} \begin{bmatrix} \zeta(x) & \phi(x) \\ \psi(x) & \mathscr{T} v_{t+1}(x) \end{bmatrix}, \quad \text{for all } x \in X, t \in [T-1],
\end{aligned}
\tag{25}
$$

where SE stands for the set of Nash (saddle-point) equilibrium values of the matrix game with two pure strategies. Then, the last iterate value function is $\zeta(\cdot)$ by construction. The rest of the business is to figure out the backward induction equation that involves the val$(\cdot)$ operator, which still relies on the calculation of $\mathscr{T}$ leveraging Monte-Carlo sampling type of methods. Following Lemma 1 the monotonicity still holds, $\{v_t(\cdot)\}_{t \in [T]}$ is decreasing, which can be interpreted as that the decision made at the outset is most valuable, as time passes, the opportunity fades. For any $t \in [T]$,

we define the two stopping times,

$$\tau_t^* = \inf\{t \le k \le T | \{v_k(X_k) = \zeta(X_k)\} \bigcup \{v_k(X_k) = \phi(X_k)\}\},$$
$$\sigma_t^* = \inf\{t \le k \le T | \{v_k(X_k) = \zeta(X_k)\} \bigcup \{v_k(X_k) = \psi(X_k)\}\}.$$

By Theorem 2, (τ_0^*, τ_0^*) is the equilibrium strategy pair, the definition of which reflects the consistency of value function computation, that is, the players' current value estimates either reach the early termination threshold or confrontational threshold.

Theorem 2. *Under ADC, the following statements hold for arbitrary initial state* $x \in \mathcal{X}$:

- *for every* $t \in [T]$, *and all* $\tau \in \mathcal{T}_t, \sigma \in \mathcal{T}_t$,

$$\mathbb{E}[H(\tau, \sigma_t^*)|\mathcal{F}_t] \le \mathbb{E}[H(\tau_t^*, \sigma_t^*)|\mathcal{F}_t] \le \mathbb{E}[H(\tau_t^*, \sigma)|\mathcal{F}_t].$$

- *the game admits a DDE strategy* (τ_0^*, σ_0^*), *at which the value function satisfies*

$$\mathcal{V}^{\tau_0^*, \sigma_0^*}(x) = \text{ess sup}_{\tau \in \mathcal{T}} \text{ess inf}_{\sigma \in \mathcal{T}} \mathcal{V}^{\tau, \sigma}(x)$$
$$= \text{ess inf}_{\sigma \in \mathcal{T}} \text{ess sup}_{\tau \in \mathcal{T}} \mathcal{V}^{\tau, \sigma}(x).$$

We omit the proof here as Theorem 2 can be seen as an extension of Theorem 1, to which the reasoning is similar. One can simply construct the sequence of value functions with a constant translation, and the results still hold.

6.2.2 Case II: Defensive Dominance

Under the Defensive Dominance Condition (DDC), the payoff functions satisfy the ordered condition $\psi \le \zeta \le \phi$ for all system states $x \in \mathcal{X}$. In this case, at any state $x \in \mathcal{X}$, the defender can bide his time during the interactions of cyber kill/defense chain, as the systematic loss after the execution of Command & Control is mitigable. Such a condition happens when the defender possesses a superior and robust position. This is also called *second-mover advantage*, that is, the defender has the incentive to wait for the opponent to end the game.

DDC corresponds to the ordered payoff condition for standard Dynkin's game, where the existence and uniqueness of a saddle point value process have been proved. The constructive sequence of (locally integrable) random variables $\{V_t\}_{t=0}^T$, in this case, is now more straightforward (as discussed in [78]), defined by

$$\begin{aligned} V_T &= \zeta(X_T), \\ V_t &= \min\{\psi(X_t), \max\{\phi(X_t), \mathbb{E}[V_{t+1}|\mathcal{F}_t]\}\}, \qquad \text{for } t = 0, \dots, T-1, \end{aligned} \tag{26}$$

with the stopping time strategies defined as

$$\bar{\tau}_t = \inf\{t \leq k \leq T | V_k = \phi(X_k)\},$$
$$\bar{\sigma}_t = \inf\{t \leq k \leq T | V_k = \psi(X_k)\}.$$

Theorem 3. *Under DDC, the following statements hold:*

i) for each $t \in [T]$, and for all $\tau \in \mathcal{T}_t$, $\sigma \in \mathcal{T}_t$,

$$V_t = \mathbb{E}[V_{\tau_t^* \wedge \sigma_t^*}|\mathcal{F}_t] = \mathbb{E}[H(\tau_t^*, \sigma_t^*)|\mathcal{F}_t], \qquad and,$$
$$\mathbb{E}[H(\tau, \sigma_t^*)|\mathcal{F}_t] \leq V_t \leq \mathbb{E}[H(\tau_t^*, \sigma)|\mathcal{F}_t].$$

ii) at every time $t \in [T]$, a pair (τ_t^*, σ_t^*) is an equilibrium point for the subgame starting at time t, and the DDE value corresponding to (τ_0^*, σ_0^*) is

$$\mathbb{E}[V_0] = \mathbb{E}[V_{\tau_0^* \wedge \sigma_0^*}] = \mathbb{E}[H(\tau_0^*, \sigma_0^*)].$$

Proof. Similar to previous results, we shall give the proof for the "$\geq$" side. First, we examine the trivial case where $t = T$. Obviously, there's no option but stop for both parties, so $G_t = \zeta(X_t) = \mathbb{E}[G_{\tau \wedge \sigma_t^*}|\mathcal{F}_t] = \mathbb{E}[H(\tau, \sigma_t^*)|\mathcal{F}_t]$ for all $\tau \in \mathcal{T}_T = \{T\}$.

Fix a $t < T$. Choose some k such that $t \leq k \leq \tau_t^* \wedge \sigma_t^*$, we have $V_k = \mathbb{E}[V_{k+1}|\mathcal{F}_k]$ by definition. Thus, $\{V_{k \wedge \tau_t^* \wedge \sigma_t^*}\}_{k=t}^T$ is a Martingale. Applying Doob's optional sampling theorem, one has

$$V_t = \mathbb{E}[V_{\tau \wedge \tau_t^* \wedge \sigma_t^*}|\mathcal{F}_t], \quad \text{for all } \tau \in \mathcal{T}_t.$$

Let $\tau = \tau_t^*$, we arrive at

$$\begin{aligned} V_t &= \mathbb{E}[V_{\tau_t^* \wedge \sigma_t^*}|\mathcal{F}_t] \\ &= \mathbb{E}[V_{\tau_t^*}\mathbb{1}_{\{\tau_t^* \leq \sigma_t^*\}} + V_{\sigma_t^*}\mathbb{1}_{\{\tau_t^* > \sigma_t^*\}}|\mathcal{F}_t] \\ &= \mathbb{E}[\phi(X_{\tau_t^*})\mathbb{1}_{\{\tau_t^* < \sigma_t^*\}} + \psi(X_{\sigma_t^*})\mathbb{1}_{\{\tau_t^* > \sigma_t^*\}} + \zeta(X_{\tau_t^*})\mathbb{1}_{\{\tau_t^* = \sigma_t^*\}}|\mathcal{F}_t] \\ &= \mathbb{E}[H(\tau_t^*, \sigma_t^*)|\mathcal{F}_t]. \end{aligned}$$

It is also obvious that when $t \leq k < \sigma_t^*$, then $V_k < \psi(X_k)$, therefore $V_k = \mathbb{E}[V_{k+1}|\mathcal{F}_k]$. This implies that $\{V_{k \wedge \sigma_t^*}\}_{k=t}^T$ is a supermartingale. Hence, $V_t \geq \mathbb{E}[V_{\tau \wedge \sigma_t^*}|\mathcal{F}_t]$ for all $\tau \in \mathcal{T}_t$,

$$\begin{aligned} \mathbb{E}[V_{\tau \wedge \sigma_t^*}|\mathcal{F}_t] &\geq \mathbb{E}[\phi(X_\tau)\mathbb{1}_{\{\tau < \sigma_t^*\}} + \psi(X_{\sigma_t^*})\mathbb{1}_{\{\tau_t^* > \sigma_t^*\}} + \zeta(X_\tau)\mathbb{1}_{\{\tau = \sigma_t^*\}}|\mathcal{F}_t] \\ &= \mathbb{E}[H(\tau, \sigma_t^*)|\mathcal{F}_t], \end{aligned}$$

since $V_k > \phi(X_k)$ and $\zeta(X_k) \leq \psi(X_k)$ for all $0 \leq k \leq T$. Claim ii) follows immediately.

Again we generalize the result to $\mathcal{E}(\mathcal{X})$, we wish to find a sequence of $\mathcal{B}(\mathbb{R})/\mathcal{G}$-measurable functions $\{v_t(\cdot)\}_{t \in [T]}$ that satisfies the following conditions (or being shifted by a constant):

$$v_T(x) = \zeta(x), \qquad \text{for all } x \in \mathcal{X},$$
$$v_t(x) = \min\{\psi(x), \max\{\phi(x), \mathcal{T}v_{t+1}(x)\}\}, \qquad \text{for all } x \in \mathcal{X}, t = [T-1],$$

and the DDE pair (τ^*, σ^*) can be defined as:

$$\tau^* = \inf\{k \in [T] | v_k(X_k) = \phi(X_k)\},$$
$$\sigma^* = \inf\{k \in [T] | v_k(X_k) = \psi(X_k)\}.$$

Theorem 4. *Under ADC, the game admits a DDE strategy pair (τ^*, σ^*), such that*

$$\mathcal{V}^{\tau^*, \sigma^*}(x) = \operatorname{ess\,sup}_{\tau \in \mathcal{T}} \operatorname{ess\,inf}_{\sigma \in \mathcal{T}} \mathcal{V}^{\tau, \sigma}(x)$$
$$= \operatorname{ess\,inf}_{\sigma \in \mathcal{T}} \operatorname{ess\,sup}_{\tau \in \mathcal{T}} \mathcal{V}^{\tau, \sigma}(x),$$

for all $x \in \mathcal{X}$.

Under DDC, the optimal strategies for the players are waiting for the equilibrium process to hit the lower/upper bound of the payoff values.

6.2.3 Decision Dominance with Information Asymmetry

In the MDW scenarios, it is crucial to recognize that both defenders and attackers operate within an environment of information asymmetry [11, 38]. This is particularly evident when considering STIX logs, as the information accessible to attackers differs from what defenders can observe. While defenders have the advantage of comprehensive logs that capture security events and indicators of compromise, attackers possess their own set of advantages stemming from their ability to exploit the gaps in the defender's knowledge. Attackers can leverage their insider information, external reconnaissance, and targeted intelligence gathering to gain insights into the defender's security measures, potential vulnerabilities, and defensive capabilities. In the meantime, the defender may have deceptive defense mechanisms that hide their tactics, techniques, and procedures (TTPs), to counteract the malicious exploitation.

To formalize the notion, we redefine $(X_t)_{t=0}^T$ as the true system state (which cannot be completely captured by the STIX logs), and let $(O_t^i)_{0 \le t \le T}(i = 1, 2)$ be the observation process for the defender $(i = 1)$ and the attacker $(i = 2)$, which jointly live in the space $(O^1 \times O^2, \mathcal{H}^1 \otimes \mathcal{H}^2)$, adapted to the filtrations $\mathbb{H}^1 = (\mathcal{H}_t^1)_{0 \le t \le T}$ and $\mathbb{H}^2 = (\mathcal{H}_t^2)_{0 \le t \le T}$. This information asymmetry enables the players to make informed decisions regarding their strategies, tactics, and the selection of attack vectors/defensive mechanisms. Therefore, defenders must not only rely on STIX logs and robust defense mechanisms but also proactively bridge the information gap by enhancing their threat intelligence capabilities, anticipating adversary behaviors, and continuously evolving their defense strategies to counter the advantages of information asymmetry in the cyber landscape.

To formally define the DDG under asymmetric information structure, we denote by $\mathcal{T}(\mathbb{H}^i)$ the set of $\mathbb{H}^i$-stopping times, $\mathcal{T}(\mathbb{H}^i) = \{0 \le \tau \le T : \{\tau(\omega) \le k\} \in \mathcal{H}_k^i \,\forall k \in [T], \forall \omega \in \Omega\}$. The decision payoffs at each stage $t \in [T]$, in this case, may depend on both O_t^i and X_t. Following the standard formalism of the Partially Observable Markov Decision Process (POMDP), we assume that the payoff functions still only depend on the true system state, which is a hidden latent variable for both

players. Instead, there exists an emission kernel $\mathbb{O} : \mathcal{X} \to \Delta(O^1 \times O^2)$ that measures the joint probability of observations made by the defender and the attacker. An illustration is shown in Figure 7
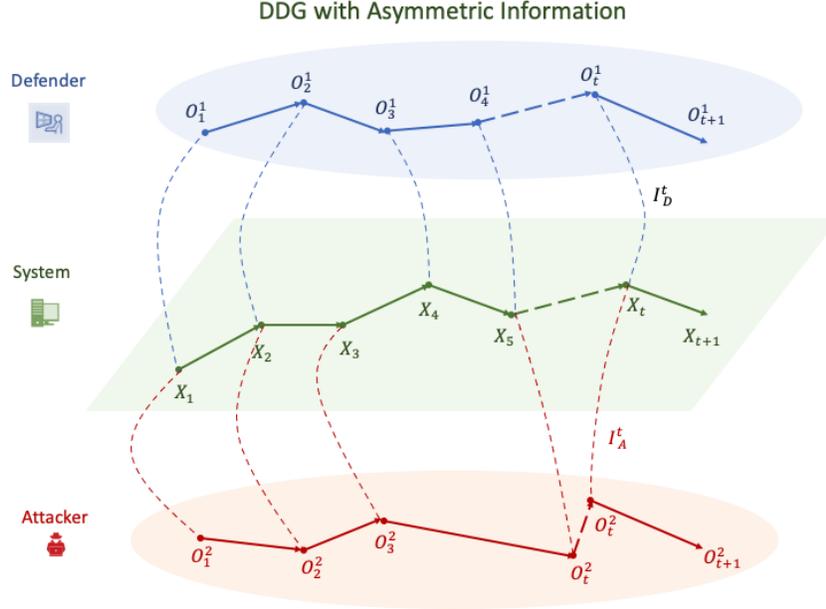
Fig. 7: An illustration of asymmetric information dynamic games defined in Definition 11. The two players have distinct partial observations for the system state X_t, denoted by (O_t^1, O_t^2). In DDG, the defender has to infer the true state to determine the stopping time strategy based on the payoff structure, which relies on credible modeling, requiring expertise in the fundamental understanding of the cyber threats

Factorization Lemma says in order to infer the true states from the partial observations, say, if O_t^1 is $\mathcal{F}_t/\mathcal{H}_t^1$-measurable, there needs to be a deterministic $\mathcal{F}_t/\mathcal{H}_t^1$-measurable map $f : \mathcal{X} \to O^1$ such that $O_t^1 = f(X_t)$, whose existence and accessibility are not always guaranteed in the cyber domain. Therefore, it is reasonable to assume that the players have their stopping time strategies restricted to $\mathcal{T}(\mathbb{H}^i)$. Definition 11 summarizes the game under asymmetrical information structure.

Definition 11 (Decision Dominance Game with Information Asymmetry).
A tuple $(\mathcal{X}, O^1 \times O^2, \mathcal{P}, \mathbb{O}, \phi, \zeta, \psi, \mathcal{T}(\mathbb{H}^1), \mathcal{T}(\mathbb{H}^2))$ encapsulates a Decision Dominance Game with Information Asymmetry (DDGIA) if it satisfies that

- There exists a hidden Markov process $(X_t)_{0 \le t \le T}$ that lives in $(X, \mathcal{G})$ with transition kernel $\mathcal{P}$, which yields observations (O_t^1, O_t^2) through emission kernel $\mathbb{O}$;
- ϕ, ζ, and ψ are payoff functions mapping from X_t to $\mathbb{R}$, $\phi, \zeta, \psi \in \mathcal{E}(X)$, which is the set of all bounded $\mathcal{B}(\mathbb{R})/\mathcal{G}$-measurable functions on $(X, \mathcal{G})$. Further, $\min(\phi, \psi) \le \zeta \le \max(\phi, \psi)$ on X;
- At each stage t, player i $(i = 1, 2)$ picks a stopping strategy from space $\mathcal{T}_t(\mathbb{H}^i) := \{t \le \tau \le T : \{\tau(\omega) \le k\} \in \mathcal{H}_k^i \ \forall k \in [T], \forall \omega \in \Omega\}$ to decide whether to stop or continue the kill/defense chain.
- At each stage the utility function of the defender is

$$H(\tau_t, \sigma_t) = \phi(X_{\tau_t})\mathbb{1}_{\{\tau_t < \sigma_t\}} + \zeta(X_{\tau_t})\mathbb{1}_{\{\tau_t = \sigma_t\}} + \psi(X_{\sigma_t})\mathbb{1}_{\{\tau_t > \sigma_t\}},$$

while the attacker attains $-H(\tau_t, \sigma_t)$.

The goal of the defender is to choose τ to maximize her utility under all possible choices of the attacker, which leads to the lower value function of DDGIA,

$$\underline{V}(x) = \operatorname{ess\,sup}_{\tau \in \mathcal{T}(\mathbb{H}^1)} \operatorname{ess\,inf}_{\sigma \in \mathcal{T}(\mathbb{H}^2)} V^{\tau, \sigma}(x). \tag{27}$$

Similarly, the goal of the attacker is to choose σ to minimize the defender's utility under all possible choices of the defender, which leads to the upper-value function,

$$\overline{V}(x) = \operatorname{ess\,inf}_{\sigma \in \mathcal{T}(\mathbb{H}^2)} \operatorname{ess\,sup}_{\tau \in \mathcal{T}(\mathbb{H}^1)} V^{\tau, \sigma}(x). \tag{28}$$

Definition 12 (DDE with Information Asymmetry). A pair of stopping time strategies $(\tau^*, \sigma^*) \in \mathcal{T}(\mathbb{H}^1) \times \mathcal{T}(\mathbb{H}^2)$ is a Decision-Dominance Equilibrium (DDE) if for all initial state $x \in X$, it satisfies the minimax condition:

$$\begin{aligned} V^{\tau^*, \sigma^*}(x) &= \operatorname{ess\,sup}_{\sigma \in \mathcal{T}(\mathbb{H}^2)} \operatorname{ess\,inf}_{\tau \in \mathcal{T}(\mathbb{H}^1)} V^{\tau, \sigma}(x) \\ &= \operatorname{ess\,inf}_{\tau \in \mathcal{T}(\mathbb{H}^2)} \operatorname{ess\,sup}_{\sigma \in \mathcal{T}(\mathbb{H}^2)} V^{\tau, \sigma}(x). \end{aligned} \tag{29}$$

We say that a DDGIA has a value if $\underline{V}(x) = \overline{V}(x)$. Note that the existence and uniqueness of the value is a non-trivial question in general, as we shall find the reasoning presented in the previous section not applicable due to the introduction of two private filtrations for both parties. In principle, the value exists if $\mathbb{H}^i$ reveal the same information from $\mathbb{F}$, in which case the conditional expectation $\mathbb{E}(\cdot|\mathbb{H}_t^i)$ can be seen equivalent with $\mathbb{E}(\cdot|\mathcal{F}_t)$, thus the players will make their decisions using the same threshold policies. This property, however, requires some special structures of the observation kernel $\mathbb{O}$, which might not hold in realistic scenarios.

6.3 Decision Dominance Zero-Trust Defense (DD-ZTD): A Case Study

In this case study, we consider an T-episodic DDG with symmetric information over the same 5G network $G = \langle V, E \rangle$ as discussed in section 5.2, where each episode t contains H ZTD steps against lateral movement. The ZTD state action variables within one episode t is $\mathbf{sa}_t = (s_t^1, a_t^1, s_t^2, a_t^2 \ldots, a_t^{H-1}, s_t^H)$, where $s_t^h = (G_t^h, L_t^h)$, $h = 1, \ldots, H$ are the authentication graphs and the visiting indicator functions at episode t, and the joint actions a_t^h are automated by the threshold-policy trust engine, which is either the Bayesian type or the Machine Learning type. Denote the STIX logs within t as $x_t \in \mathcal{X}$, which includes but is not limited to the events of 5G network exposure, slicing control, session management; the threat actor characterizations such as suspected user intentions and handling guidance. The Markovian state at episode t is a composition of both historical ZTD state action variables and the STIX logs gathered before episode t, i.e., $X_t := (\mathbf{sa}_{1:t-1}, x_{t-1})$.

During the cyber kill/defense chain interaction, at the beginning of each episode, the defender can choose to completely cut off the chain before episode t starts by isolating the networks, restarting the services, resetting all the credentials, patching and hardening the security configurations, and then restoring and resuming the operations. The cost of the defender's cutting-off strategy is $C(\cdot) : \mathcal{X} \to \mathbb{R}$, which only depends on the cyber threat information. Similarly, the attacker can choose to take action early by exploiting Zero-Day vulnerabilities, evading intrusion detection systems, and implementing stealthy command and control at an early stage of the cyber kill chain. Again we let the exploitation loss be $\ell(\cdot) : \mathcal{X} \to \mathbb{R}$, which completely depends on the cyber threat characterization of episode t. Now we are ready to define the three payoff functions in our DDG framework.

The early termination payoff, confrontation payoff, and late termination payoff functions can be defined as

$$\phi(X_t) = -\mathbb{E}\Big[\sum_{k=1}^{t-1}\sum_{h=1}^{H} u_D(s_k^h, a_k^h)\Big] - C(x_{t-1}),$$

$$\zeta(X_t) = -\mathbb{E}\Big[\sum_{k=1}^{t-1}\sum_{h=1}^{H} u_D(s_k^h, a_k^h)\Big] - C(x_{t-1}) - \ell(x_{t-1}), \qquad (30)$$

$$\psi(X_t) = -\mathbb{E}\Big[\sum_{k=1}^{t}\sum_{h=1}^{H} u_D(s_k^h, a_k^h)\Big] - \ell(x_{t-1}),$$

where the expectation $\mathbb{E}[\sum_h^H u_D(s_t^h, a_t^h)]$ is taken conditioned on $\mathbf{sa}_{1:t-1}$. The interpretation is that when the defender chooses to shut down and restore the services, the ZTD stops for that episode, while if the attacker chooses to exploit early, the ZTD mechanism is still active.

One can easily verify that when both ℓ and C are positive and the expected ZTD cost within every episode t satisfies $\mathbb{E}[\sum_h^H u_D(s_t^h, a_t^h)] > C(x_{t-1})$ the DDG satisfies DDC.

7 Conclusion

This chapter develops a game-theoretic framework for the decision-dominant zero-trust defense of 5G networks in the face of advanced persistent threats that utilize a cyber kill chain to disrupt the network operation. The advanced features of 5G networks, despite their contributions to multi-domain integration, bring a larger attack surface and render the network system vulnerable in the presence of advanced persistent threats (APT) and other malicious attacks. The combination of vulnerabilities in APT, supply chains of 5G equipment, and network slicing, along with others, can be exploited by an APT attacker to create a cyber kill chain consisting of reconnaissance, planning, execution, and exploration.

To outmaneuver the malicious attacker and thwart the kill chain, this chapter proposes a decision-dominant zero-trust defense (DD-ZTD) framework, a proactive defense mechanism enabling the defender to make timely and effective decisions with incomplete information regarding the situation and disrupt the kill chain before its completion. Two pillars of DD-ZTD are game-theoretic zero-trust defense built upon asymmetric information Markov games (AIMG) and decision-dominance defense characterized by Dykin's stopping-time games. With the AIMG's expressive power on information structures in cyber defense, ZTD employs a variety of trust engines to evaluate entities' trustworthiness with limited partial observations, which is then fed into the access policy powered by equilibrium thinking that anticipates the attacker's response. We further present an end-to-end ZTD facilitated by recent machine learning advancements with data-driven trust evaluation and explainable and generalizable policy learning.

While the proposed ZTD offers a set of fruitful tools to quantitatively analyze trustworthiness under information asymmetry, the networked entities still face multistage persistent cyber threats that call for rapid response from the defender. To outpace the attacker's kill chain, decision-dominance defense (D^3), mathematically treating interactions of cyber defense/kill chain as a stopping-time game, aims to take the decisive move to cut off the kill chain before the attack materializes. The essence of D^3 is the timing of the cutting-off, which is determined by the equilibrium of the game with anticipation of the attacker's strategic move. The resulting DD-ZTD, as an organic integration of the two game-theoretic defense mechanisms, displays great potential in combating sophisticated adversaries, which we articulate using a case study in 5G network defense.

References

[1] Headquarters, Department of the Army (2022) FM 3-0, Operations. `https://usacac.army.mil/node/3048`, accessed 2023-07-05

[2] Department of Defense (2018) Summary of the 2018 national defense strategy. `https://www.spoc.spaceforce.mil/About-Us/Fact-Sheets/Display/Article/2381348/advanced-extremely-high-frequency-system-aehf`, accessed 2023-07-05

[3] Space Operations Command (SPOC) (2021) Advanced extremely high frequency system (aehf). `https://www.spoc.spaceforce.mil/About-Us/Fact-Sheets/Display/Article/2381348/advanced-extremely-high-frequency-system-aehf`, accessed 2023-07-05

[4] Lockheed Martin (2023) Indago UAV. `https://www.lockheedmartin.com/en-us/products/indago-vtol-uav.html`, accessed 2023-07-05

[5] Huang L, Zhu Q (2022) Radams: Resilient and adaptive alert and attention management strategy against informational denial-of-service (idos) attacks. Computers & Security 121:102844

[6] Wijethilaka S, Liyanage M (2021) Survey on network slicing for internet of things realization in 5g networks. IEEE Communications Surveys & Tutorials 23(2):957–994, DOI 10.1109/COMST.2021.3067807

[7] Xiao Y, Jia Y, Liu C, Cheng X, Yu J, Lv W (2019) Edge computing security: State of the art and challenges. Proceedings of the IEEE 107(8):1608–1631, DOI 10.1109/JPROC.2019.2918437

[8] Zhu Q, Rass S (2018) On multi-phase and multi-stage game-theoretic modeling of advanced persistent threats. IEEE Access 6:13958–13971

[9] Huang L, Zhu Q (2020) A dynamic games approach to proactive defense strategies against advanced persistent threats in cyber-physical systems. Computers & Security 89:101660

[10] Rass S, Zhu Q (2016) Gadapt: a sequential game-theoretic framework for designing defense-in-depth strategies against advanced persistent threats. In: International conference on decision and game theory for security, Springer, pp 314–326

[11] Huang L, Zhu Q (2019) Dynamic bayesian games for adversarial and defensive cyber deception. Autonomous Cyber Deception: Reasoning, Adaptive Planning, and Evaluation of HoneyThings pp 75–97

[12] Rose S, Borchert O, Mitchell S, Connelly S (2020) Zero trust architecture. Tech. rep., National Institute of Standards and Technology

[13] Osborn K (2018) "first look, first shot, first kill": How the f-22 raptor could fly until 2060. URL `https://nationalinterest.org/blog/buzz/first-look-first-shot-first-kill-how-f-22-raptor-could-fly-until-2060-35937`

[14] Gady FS, Stronell A (2020) Cyber capabilities and multi-domain operations in future high-intensity warfare in 2030. Cyber Threats and NATO 2030: Horizon Scanning and Analysis pp 151–176

[15] Mallik RK, Scholtz RA, Papavassilopoulos GP (2000) Analysis of an on-off jamming situation as a dynamic game. IEEE Transactions on Communications 48(8):1360–1373

[16] Mukherjee A, Swindlehurst AL (2012) Jamming games in the mimo wiretap channel with an active eavesdropper. IEEE Transactions on Signal Processing 61(1):82–91

[17] Sayin MO, Hosseini H, Poovendran R, Başar T (2018) A game theoretical framework for inter-process adversarial intervention detection. In: International Conference on Decision and Game Theory for Security, Springer, pp 486–507

[18] Chen J, Touati C, Zhu Q (2019) Optimal secure two-layer iot network design. IEEE Transactions on Control of Network Systems pp 1–1, DOI 10.1109/TCNS.2019.2906893

[19] Pawlick J, Farhang S, Zhu Q (2015) Flip the cloud: Cyber-physical signaling games in the presence of advanced persistent threats. In: International Conference on Decision and Game Theory for Security, Springer, pp 289–308

[20] Pawlick J, Zhu Q (2017) Strategic trust in cloud-enabled cyber-physical systems with an application to glucose control. IEEE Transactions on Information Forensics and Security 12(12):2906–2919

[21] Huang L, Chen J, Zhu Q (2017) A large-scale markov game approach to dynamic protection of interdependent infrastructure networks. In: International Conference on Decision and Game Theory for Security, Springer, pp 357–376

[22] Chen J, Zhu Q (2022) A cross-layer design approach to strategic cyber defense and robust switching control of cyber-physical wind energy systems. IEEE Transactions on Automation Science and Engineering 20(1):624–635

[23] Chen J, Zhu Q (2019) A game-and decision-theoretic approach to resilient interdependent network analysis and design. Springer

[24] Chen J, Zhu Q (2016) A game-theoretic framework for resilient and distributed generation control of renewable energies in microgrids. IEEE Transactions on Smart Grid 8(1):285–295

[25] Chen J, Zhu Q (2019) A games-in-games approach to mosaic command and control design of dynamic network-of-networks for secure and resilient multi-domain operations. In: Chen G, Pham KD (eds) Sensors and Systems for Space Applications XII, International Society for Optics and Photonics, SPIE, vol 11017, pp 189 – 195, DOI 10.1117/12.2526677, URL https://doi.org/10.1117/12.2526677

[26] Chen J, Zhu Q (2020) Control of multilayer mobile autonomous systems in adversarial environments: A games-in-games approach. IEEE Transactions on Control of Network Systems 7(3):1056–1068, DOI 10.1109/TCNS.2019.2962316

[27] Chen J, Zhu Q (2016) Resilient and decentralized control of multi-level cooperative mobile networks to maintain connectivity under adversarial environment. In: IEEE Conference on Decision and Control (CDC), pp 5183–5188

[28] Zhu Q, Rass S, Dieber B, Vilches VM, et al. (2021) Cybersecurity in robotics: Challenges, quantitative modeling, and practice. Foundations and Trends® in Robotics 9(1):1–129

[29] Kieras T, Farooq MJ, Zhu Q (2020) Riots: Risk analysis of IoT supply chain threats. In: 2020 IEEE 6th World Forum on Internet of Things (WF-IoT), IEEE, pp 1–6

[30] Ge Y, Zhu Q (2022) Accountability and insurance in IoT supply chain. arXiv preprint arXiv:220111855 URL `https://doi.org/10.48550/arXiv.2201.11855`

[31] Kieras T, Farooq J, Zhu Q (2022) IoT supply chain security risk analysis and mitigation: Modeling, computations, and software tools. Springer

[32] Pan Y, Zhu Q (2022) On poisoned wardrop equilibrium in congestion games. In: International Conference on Decision and Game Theory for Security, Springer, pp 191–211

[33] Pan Y, Li T, Zhu Q (2023) On the resilience of traffic networks under non-equilibrium learning. In: 2023 American Control Conference (ACC), IEEE, pp 3484–3489

[34] Pan Y, Li T, Zhu Q (2023) Is stochastic mirror descent vulnerable to adversarial delay attacks? a traffic assignment resilience study. arXiv preprint arXiv:230401161 URL `https://doi.org/10.48550/arXiv.2304.01161`

[35] Zheng J, Castañón DA (2012) Dynamic network interdiction games with imperfect information and deception. In: 2012 IEEE 51st IEEE Conference on Decision and Control (CDC), IEEE, pp 7758–7763

[36] Zhu Q, Clark A, Poovendran R, Başar T (2012) Deceptive routing games. In: 2012 IEEE 51st IEEE Conference on Decision and Control (CDC), IEEE, pp 2704–2711

[37] Zhuang J, Bier VM, Alagoz O (2010) Modeling secrecy and deception in a multiple-period attacker–defender signaling game. European Journal of Operational Research 203(2):409–418

[38] Pawlick J, Colbert E, Zhu Q (2019) A game-theoretic taxonomy and survey of defensive deception for cybersecurity and privacy. ACM Computing Surveys (CSUR) 52(4):82

[39] Zhu Q, Başar T (2013) Game-theoretic approach to feedback-driven multi-stage moving target defense. In: Decision and Game Theory for Security, Springer, pp 246–263

[40] Jajodia S, Ghosh AK, Swarup V, Wang C, Wang XS (2011) Moving target defense: creating asymmetric uncertainty for cyber threats, vol 54. Springer Science & Business Media

[41] Huang L, Zhu Q (2021) Combating informational denial-of-service (idos) attacks: Modeling and mitigation of attentional human vulnerability. In: Decision and Game Theory for Security: 12th International Conference, GameSec 2021, Virtual Event, October 25–27, 2021, Proceedings, Springer, pp 314–333

[42] Huang L, Zhu Q (2023) Cognitive security: A system-scientific approach. Springer Nature

[43] Liao HJ, Richard Lin CH, Lin YC, Tung KY (2013) Intrusion detection system: A comprehensive review. Journal of Network and Computer Applications 36(1):16–24, DOI https://doi.org/10.1016/j.jnca.2012.09.004, URL `https://www.sciencedirect.com/science/article/pii/S1084804512001944`

[44] Bhatt S, Manadhata PK, Zomlot L (2014) The operational role of security information and event management systems. IEEE Security & Privacy 12(5):35–41, DOI 10.1109/msp.2014.103

[45] Li T, Zhao Y, Zhu Q (2022) The role of information structures in game-theoretic multi-agent learning. Annual Reviews in Control 53:296–314, DOI 10.1016/j.arcontrol.2022.03.003

[46] Li T, Zhu Q (2022) Commitment with signaling under double-sided information asymmetry. arXiv preprint arXiv:221211446 URL `https://doi.org/10.48550/arXiv.2212.11446`

[47] Ge Y, Li T, Zhu Q (2023) Scenario-agnostic zero-trust defense with explainable threshold policy: A meta-learning approach. In: IEEE INFOCOM 2023 - IEEE Conference on Computer Communications Workshops (INFOCOM WKSHPS), pp 1–6, DOI 10.1109/INFOCOMWKSHPS57453.2023.10225816

[48] Ometov A, Bezzateev S, Mäkitalo N, Andreev S, Mikkonen T, Koucheryavy Y (2018) Multi-factor authentication: A survey. Cryptography 2(1):1

[49] OpenAI (2023) Gpt-4 technical report. arXiv preprint arXiv:230308774 URL `https://doi.org/10.48550/arXiv.2303.08774`

[50] Vaswani A, Shazeer N, Parmar N, Uszkoreit J, Jones L, Gomez AN, Kaiser Lu, Polosukhin I (2017) Attention is all you need. In: Guyon I, Luxburg UV, Bengio S, Wallach H, Fergus R, Vishwanathan S, Garnett R (eds) Advances in Neural Information Processing Systems, Curran Associates, Inc., vol 30, URL `https://proceedings.neurips.cc/paper_files/paper/2017/file/3f5ee243547dee91f\bd053c1c4a845aa-Paper.pdf`

[51] Kingma DP, Welling M (2014) Auto-encoding variational bayes. In: 2nd International Conference on Learning Representations, ICLR 2014, Banff, AB, Canada, April 14-16, 2014, Conference Track Proceedings, `http://arxiv.org/abs/1312.6114v10`

[52] Paisley J, Blei DM, Jordan MI (2012) Variational bayesian inference with stochastic search. In: Proceedings of the 29th International Coference on International Conference on Machine Learning, Omnipress, Madison, WI, USA, ICML'12, p 1363–1370

[53] Nash J (1951) Non-cooperative games. The Annals of Mathematics 54(2):286–295, DOI 10.2307/1969529

[54] Fudenberg D, Tirole J (1991) Game Theory. MIT Press, Cambridge, MA

[55] Li T, Zhu Q (2023) On the price of transparency: A comparison between overt persuasion and covert signaling. arXiv preprint arXiv:230400096 URL `https://doi.org/10.48550/arXiv.2304.00096`

[56] Li T, Zhu Q (2019) On convergence rate of adaptive multiscale value function approximation for reinforcement learning. 2019 IEEE 29th International Workshop on Machine Learning for Signal Processing (MLSP) pp 1–6, DOI 10.1109/mlsp.2019.8918816

[57] Li T, Peng G, Zhu Q (2021) Blackwell online learning for markov decision processes. 2021 55th Annual Conference on Information Sciences and Systems (CISS) 00:1–6, DOI 10.1109/ciss50987.2021.9400319

[58] Sutton RS, McAllester DA, Singh SP, Mansour Y (2000) Policy gradient methods for reinforcement learning with function approximation. In: Advances in Neural Information Processing Systems 12, MIT press, pp 1057—1063, URL `http://papers.nips.cc/paper/1713-policy-gradient-methods-for-reinforcement-learning-with-function-approximation.pdf`

[59] Bannon J, Windsor B, Song W, Li T (2020) Causality and batch reinforcement learning: Complementary approaches to planning in unknown domains. arXiv preprint arXiv:200602579 URL `https://doi.org/10.48550/arXiv.2006.02579`

[60] Puterman ML (1994) Markov decision processes: Discrete stochastic dynamic programming, 1st edn. John Wiley & Sons, Inc., USA

[61] Hu J, Wellman MP (2003) Nash q-learning for general-sum stochastic games. Journal of machine learning research 4(Nov):1039—1069

[62] Hammar K, Stadler R (2023) Digital twins for security automation. In: NOMS 2023-2023 IEEE/IFIP Network Operations and Management Symposium, pp 1–6, DOI 10.1109/NOMS56928.2023.10154288

[63] Li T, Lei H, Zhu Q (2022) Sampling attacks on meta reinforcement learning: A minimax formulation and complexity analysis. arXiv preprint arXiv:220800081 URL `https://doi.org/10.48550/arXiv.2208.00081`

[64] Dazeley R, Vamplew P, Cruz F (2023) Explainable reinforcement learning for broad-xai: A conceptual framework and survey. Neural Comput Appl 35(23):16893–16916, DOI 10.1007/s00521-023-08423-1, URL `https://doi.org/10.1007/s00521-023-08423-1`

[65] Ge Y, Zhu Q (2022) Trust threshold policy for explainable and adaptive zero-trust defense in enterprise networks. In: 2022 IEEE Conference on Communications and Network Security (CNS), pp 359–364, DOI 10.1109/CNS56114.2022.9947263

[66] Hospedales TM, Antoniou A, Micaelli P, Storkey AJ (2021) Meta-learning in neural networks: A survey. IEEE Transactions on Pattern Analysis and Machine Intelligence PP(99):1–1, DOI 10.1109/tpami.2021.3079209

[67] Pan Y, Li T, Li H, Xu T, Zheng Z, Zhu Q (2023) A first order meta stackelberg method for robust federated learning. arXiv preprint arXiv:230613800 URL `https://doi.org/10.48550/arXiv.2306.13800`

[68] Vapnik V (1999) The nature of statistical learning theory. Springer science & business media

[69] Liu S, Li T, Zhu Q (2023) Game-theoretic distributed empirical risk minimization with strategic network design. IEEE Transactions on Signal and Information Processing over Networks 9:542–556, DOI 10.1109/TSIPN.2023.3306106

[70] Strom BE, Applebaum A, Miller DP, Nickels KC, Pennington AG, Thomas CB (2018) Mitre att&ck: Design and philosophy. In: Technical report, The MITRE Corporation

[71] Hochreiter SY (2001) Learning to learn using gradient descent. Lecture Notes in Computer Science pp 87–94, DOI 10.1007/3-540-44668-0_13

[72] Li Z, Zhou F, Chen F, Li H (2017) Meta-sgd: Learning to learn quickly for few-shot learning. arXiv preprint arXiv: 170709835 URL https://doi.org/10.48550/arXiv.1707.09835

[73] Yadav T, Rao AM (2015) Technical aspects of cyber kill chain. In: Security in Computing and Communications: Third International Symposium, SSCC 2015, Kochi, India, August 10-13, 2015. Proceedings 3, Springer, pp 438–452

[74] Khan MS, Siddiqui S, Ferens K (2018) A cognitive and concurrent cyber kill chain model. Computer and Network Security Essentials pp 585–602

[75] Huang L, Zhu Q (2019) Adaptive honeypot engagement through reinforcement learning of semi-markov decision processes. In: Decision and Game Theory for Security: 10th International Conference, GameSec 2019, Stockholm, Sweden, October 30–November 1, 2019, Proceedings 10, Springer, pp 196–216

[76] Heckman KE, Stech FJ, Schmoker BS, Thomas RK (2015) Denial and deception in cyber defense. Computer 48(4):36–44

[77] Gore R, Padilla J, Diallo S (2017) Markov chain modeling of cyber threats. The Journal of Defense Modeling and Simulation 14(3):233–244

[78] Kingman J (1976) Discrete-parameter martingales