1540 TechHorizon

United Nations Security Council Resolution 1540 and Advances in Science and Technology

Special edition of the 1540 Tech-Horizon initiative



WARWICK THE UNIVERSITY OF WARWICK | KING'S College LONDON | UCL | Cranfield University | UNODA UNITED NATIONS OFFICE FOR DISARMAMENT AFFAIRS

Funded by the EU

PUBLICATION INFORMATION

United Nations Security Council Resolution 1540 and Advances in Science and Technology

Published on 11th August, 2026,
at University of Warwick, Coventry,
CV4 7AL, UK

Disclaimer

The view expressed in the publication, including the use of certain terminology, are those of the authors and do not necessarily reflect the official positions of the University of Warwick, its partner institutions, the United Nations, or the European Union.



Foreword

As an international community, we are witnessing one of the most profound technological evolutions in history. Technologies ranging from artificial intelligence to synthetic biology are transforming societies around the world. These technologies have immense potential for good, but if misused, they could also pose significant challenges to international peace and security.

The acquisition by non-State actors of nuclear, chemical and biological weapons and their means of delivery has long been a serious concern for Member States of the United Nations. At a time of geostrategic instability, scientific and technological developments may increase this threat by lowering barriers to the acquisition of weapons of mass destruction, accelerating the diffusion of sensitive knowledge, and creating new opportunities for the misuse of materials, technologies and expertise. These evolving risks reinforce the importance of the effective implementation of United Nations Security Council resolution 1540 (2004).

Resolution 1540 (2004) recognizes the value of engaging a broad range of stakeholders, including academia, in preventing the proliferation of weapons of mass destruction to and by non-State actors. Through research, innovation and education, the academic community helps raise awareness of proliferation risks, promote responsible science and inform dialogue on emerging challenges.

This Special Issue is a timely contribution to understanding the implications of scientific and technological developments for the implementation of resolution 1540 (2004). It is also an outcome of the TechHorizon initiative, launched by a consortium of universities across the United Kingdom in partnership with the United Nations Office for Disarmament Affairs (UNODA). By bringing together scientific, academic and policy communities, the initiative advances dialogue and understanding of emerging technologies in support of international peace and security.

I wish to express my sincere appreciation to the University of Warwick, the contributing authors and all partners involved in this initiative for their dedication and expertise. I also thank the European Union for its generous support, which made this publication possible.

I hope this Special Issue will inspire continued research, dialogue and collaboration among scholars, policymakers and practitioners. By strengthening responsible stewardship of scientific and technological progress, we can enhance our collective capacity to address evolving proliferation challenges and reinforce international peace and security.

Izumi Nakamitsu
Under-Secretary-General and
High Representative for Disarmament Affairs
United Nations

SPECIAL EDITION

Contents I

1 **Advanced Technologies and the Implementation of United Nations Security Council Resolution 1540** 4
Xiaohui Wu

2 **UN Security Council Resolution 1540 (2004) at the Final Frontier: Development of National Space Regulation Frameworks and UNSCR 1540** 9
Andrew Horton

3 **Artificial Intelligence as a Tool for Preventing the Proliferation of Weapons of Mass Destruction: A Civilian Perspective** 24
Sai Gu

4 **AI-Enabled Export and Border-Control Awareness under UNSCR 1540: A Narrative Overview through Goods, Data and Routes** 41
Yu Guan; Sai Gu

5 **Aerial Drones and "Means of Delivery" in the Context of UN Security Council Resolution 1540 (2004)** 57
Matthew J. F. Healy; Andrew Horton

6 **AI and Synthetic Biology Convergence: Examining approaches to corporate self-governance and implications for Security Council Resolution 1540** 69
Hassan Elbahtimy

SPECIAL EDITION

Contents II

7 **Knowledge Proliferation, UNSCR 1540, and Research Security: Governing Emerging Technologies** — Chi Zhang — 84

8 **From Material Control to Capability Formation: Proliferation-Sensitive Knowledge and UNSCR 1540 Governance** — Nanda Kishore — 96

9 **From Classroom to Commercialisation: Biomedical Engineering Capability Formation and UNSCR 1540 Governance** — Spyros Kamins — 112

10 **Cryptocurrencies and Resolution 1540 (2004)** — Jonathan Brewer — 135

11 **Making Emerging Risks Visible under UNSCR 1540: Academic Knowledge Networks and the Conditions for Distributed Anticipatory Infrastructure** — Konstantinos Papadikis — 143

ARTICLE 01

Advanced Technologies and the Implementation of United Nations Security Council Resolution 1540

Xiaohui Wu
Head of the 1540 Support Unit
UNODA



[This article is intended to provide an introduction to the growing attention being paid to advanced technologies in the context of United Nations Security Council resolution 1540 (2004) and related deliberations within the United Nations. It aims to offer a factual and informative overview of how emerging and evolving technologies may inform discussions concerning the prevention of the proliferation of weapons of mass destruction (WMD) to and by non-State actors under the framework of resolution 1540, including both the challenges they may pose and the opportunities they may offer for effective implementation of the resolution.]

Introduction

United Nations Security Council resolution 1540 (2004) was adopted to prevent non-State actors from acquiring, developing, manufacturing, possessing, transporting, transferring or using nuclear, chemical and biological weapons and their means of delivery. The resolution establishes binding obligations on all States to implement appropriate legal, regulatory, enforcement, export-control and other relevant measures to prevent such proliferation risks.

When resolution 1540 (2004) was adopted, many of today's most transformative technologies—including advanced artificial intelligence, synthetic biology, additive manufacturing, cloud computing, autonomous systems and sophisticated cyber capabilities—either did not exist or were at a nascent stage of development. Over the more than two decades since the resolution's adoption, rapid scientific and technological advances have significantly altered the proliferation landscape. While resolution 1540 remains technology-neutral and its fundamental objectives remain unchanged, these developments have heightened the need to ensure effective implementation of its obligations in a rapidly evolving scientific and technological environment.

References in subsequent United Nations Security Council resolutions of resolution 1540

United Nations Security Council Resolution 2325 (2016)

Resolution 2325 (2016) was adopted following the first comprehensive review of the status of implementation of resolution 1540 (2004). The review found that, while significant progress had been made, implementation remained uneven, while the proliferation landscape continued to evolve. In this context, the Security Council recognized the potential implications of advances in science and technology for the implementation of resolution 1540 and encouraged both Member States and the 1540 Committee to take into account emerging challenges associated with such developments.

Resolution 2325 (2016) was the first Security Council resolution to explicitly acknowledge the implications of scientific and technological change for implementation efforts. In operative paragraph 7, the Security Council:

> Calls upon States to take into account developments on the evolving nature of risk of proliferation and rapid advances in science and technology in their implementation of resolution 1540 (2004).

The resolution further recognized the need for the 1540 Committee itself to remain attentive to these developments. Operative paragraph 8:

> Requests the 1540 Committee to take note in its work, where relevant, of the continually evolving nature of the risks of proliferation, including the use by non-State actors of rapid advances in science, technology and international commerce for proliferation purposes, in the context of the implementation of resolution 1540 (2004).

These provisions reflected a growing recognition that proliferation risks are dynamic and that the effective implementation of resolution 1540 requires continued attention to developments in science and technology.

Security Council Resolution 2663 (2022) was adopted following the second comprehensive review of the status of implementation of resolution 1540 and renewed the mandate of the 1540 Committee for a further ten years, through 2032. The review noted both the progress achieved by States and the continuing challenges associated with the changing proliferation environment, including developments in science and technology and the increasingly complex methods available to non-State actors.

Reflecting these concerns, the Security Council expressed in the preamble to the resolution its continuing concern regarding:

> The threat of terrorism and the risk that non-State actors may acquire, develop, traffic in or use nuclear, chemical, and biological weapons and their means of delivery, and related materials, including by using the rapid advances in science, technology and international commerce to that end.

The Council then reaffirmed the language first introduced in resolution 2325. It:

> Again calls upon States to take into account developments on the evolving nature of risk of proliferation and rapid advances in science and technology in their implementation of resolution 1540 (2004);

and

> Reiterates its request to the 1540 Committee to take note in its work, where relevant, of the continually evolving nature of the risks of proliferation, including the use by non-State actors of rapid advances in science, technology and international commerce for proliferation purposes, in the context of the implementation of resolution 1540 (2004).

The language contained in resolutions 2325 (2016) and 2663 (2022) does not identify any specific technologies that pose particular threats. Rather, it recognizes that scientific and technological progress may have implications for the proliferation landscape and may be relevant to national implementation efforts and the work of the 1540 Committee. This approach is consistent with the broader objective of resolution 1540: ensuring the continued effectiveness of legal, regulatory, enforcement, export-control and other relevant measures to prevent non-State actors from acquiring weapons of mass destruction and related materials.

Relevant Discussions at the United Nations Security Council

In the context of continuing advances in science and technology over the past more than two decades, discussions within the 1540 community have increasingly considered the implications of emerging technologies such as artificial intelligence, synthetic biology, cybersecurity, additive manufacturing and other dual-use technologies relevant to non-proliferation efforts. While the technologies under consideration have evolved, these discussions reflect a longstanding awareness that technological change can affect proliferation risks and that the effective implementation of resolution 1540 is supported by continued attention to developments in science and technology.

At various meetings related to resolution 1540, including open briefings of the 1540 Committee, annual briefings by the 1540 Committee Chair to the Security Council, and joint briefings by the Security Council's subsidiary bodies (the 1540 Committee, the 1267/1989/2253 ISIL (Da'esh) and Al-Qaida Sanctions Committee, and the Counter-Terrorism Committee), Member States have increasingly drawn attention to the implications of rapid scientific and technological developments for the implementation of resolution 1540 and efforts to prevent the proliferation of weapons of mass destruction to non-State actors.

Many Member States highlighted the proliferation risks associated with emerging technologies, including artificial intelligence, synthetic biology, gene editing, additive manufacturing (3D printing), cyber capabilities, autonomous systems, unmanned technologies and other rapidly evolving technologies. Several delegations noted that advances in science and technology, combined with the increasing accessibility of dual-use technologies and capabilities and the wider diffusion of technical knowledge, are changing the proliferation landscape and lowering barriers to access and potential misuse by non-State actors.

A recurring concern was that terrorist groups and other non-State actors could exploit emerging technologies to acquire weapons of mass destruction, related materials, expertise or capabilities. In particular, concerns were raised regarding the potential misuse of artificial intelligence, synthetic biology and cyber tools, as well as the increasing transfer of intangible knowledge and the risks posed by cyber vulnerabilities affecting nuclear, chemical and biological facilities.

Many Member States stressed that proliferation risks are evolving alongside scientific and technological developments and that implementation of resolution 1540 must adapt accordingly. Delegations emphasized the importance of vigilance, proactive measures and international cooperation to address emerging risks, including through appropriate governance and oversight of emerging technologies and strengthened efforts to prevent their misuse by non-State actors.

Several interventions highlighted the role of the 1540 Committee in helping Member States remain informed about emerging proliferation threats and supporting effective implementation of the resolution in a changing technological environment. Some Member States also emphasized the importance of ensuring that existing tools, methodologies and assessment mechanisms remain fit for purpose and adequately reflect evolving scientific and technological developments. Support was also expressed for enhanced coordination among relevant Security Council subsidiary bodies in addressing challenges related to non-proliferation, counter-terrorism, and emerging technologies.

Numerous Member States have consistently emphasized that rapid advances in science and technology are creating new and increasingly complex proliferation risks, reinforcing the need for continued vigilance, international cooperation and effective implementation of resolution 1540 to prevent the acquisition of weapons of mass destruction by non-State actors.

Leveraging Advanced Technologies for Resolution 1540 Implementation

Scientific and technological developments are influencing not only the risks addressed by resolution 1540, but also the tools available to States for implementing its obligations. Technologies used in customs administration, border management, strategic trade controls, law enforcement and supply-chain management may enhance national efforts to prevent non-State actors from acquiring weapons of mass destruction, their means of delivery and related materials.

Technological innovation has also created new opportunities for international cooperation and assistance. Digital communication platforms, virtual training tools and other online resources have broadened access to technical expertise and facilitated exchanges among States, international organizations, industry and academia. Such developments could support implementation efforts and contribute to capacity-building in areas relevant to resolution 1540.

Reflecting the growing attention being paid to developments in science and technology in the context of resolution 1540, UNODA's recently updated online course on resolution 1540 addresses issues associated with advanced and emerging technologies, including examples of technologies with dual-use characteristics and their potential implications for implementation. The inclusion of this topic reflects efforts to ensure that resources to support capacity-building and implementation remain responsive to the evolving technological landscape and emerging proliferation challenges.

As the international community adapts to a rapidly evolving technological environment, consideration of both the risks and potential benefits associated with technological innovation may help ensure that new capabilities are effectively integrated into national implementation measures. In this regard, advances in science and technology may serve not only as sources of new proliferation challenges, but also as tools that can support the effective implementation of resolution 1540.

Conclusion

The growing attention to advanced technologies in discussions related to resolution 1540 reflects the continued relevance of the resolution in addressing evolving proliferation risks.

As scientific and technological innovation continues to evolve, exchanges among Member States, international organizations, experts and other stakeholders provide opportunities to share perspectives on emerging challenges. Ongoing dialogue, knowledge-sharing and capacity-building may support the effective implementation of resolution 1540 and broader international efforts aimed at preventing the proliferation of nuclear, chemical and biological weapons and their means of delivery to and by non-State actors.

(*Xiaohui Wu is Head of the UNSCR 1540 Support Unit, Weapons of Mass Destruction Branch, United Nations Office for Disarmament Affairs*)

ARTICLE 02

UN Security Council Resolution 1540 (2004) at the Final Frontier: Development of National Space Regulation Frameworks and UNSCR 1540

Andrew Horton

Visiting Senior Research Fellow

King's Collage London, Centre for Science & Security Studies (CSSS), Department of War Studies



Keywords

Hague Code of Conduct, hypersonic, means of delivery, missiles, MTCR, remote sensing, space launch, space technology, weapons of mass destruction, X as a Service (XaaS)

Abstract

UN Security Council resolution (UNSCR) 1540 (2004) places specific international legally binding obligations on states in respect to preventing the proliferation of nuclear, chemical and biological weapons, and 'their means of delivery' to, and by, non-state actors. Under the current international non-proliferation architecture, there are several legally binding instruments that are specifically concerned with the proliferation of the weapons themselves. As these regimes often form the framework of national legislation relevant to the implementation of UNSCR 1540 they are seen as complementary to the resolution. However, in the case of means of delivery, there is a lack of any equivalent international legally binding instrument around which states can build a national regulatory framework. There are some limited complementarities between the UNSCR 1540, and non-binding arrangements such as the Hague Code of Conduct against Ballistic Missile Proliferation (HCoC) and the Missile Technology Control Regime (MTCR), (although in the case of MTCR this is generally limited to export controls). This

article aims to explore both the complementarities within the existing, if somewhat limited, non-proliferation mechanisms, but also to consider the potential for wider complementarities with legally binding instruments, and non-binding mechanisms, outside of the traditional non-proliferation architecture.

The article highlights the potential usefulness of using national space regulation, especially when it comes to non-state commercial space actors as the holders of sensitive launch vehicle technologies relevant to ballistic and other advanced missile systems. Whilst the focus of the existing limited complementarities between UNSCR 1540 implementation and national space regulatory frameworks, has been the risk of the proliferation of sensitive missile related technology, there is a potential growing area of concern that may be addressed through such domestic space regulations, and that is the provision of space-based services that may contribute to the planning, deployment and ultimately the use of WMD by violent non-state actors. Therefore, ideally states should consider when developing a national space legal framework, not just the regulation of missile related space launch activities, but also the potential regulation of space-based services, which may inadvertently facilitate the use of nuclear, chemical and biological weapons, and their means of delivery.

Introduction

UN Security Council resolution (UNSCR) 1540 (2004) places legally binding obligations on all UN member states in respect to the proliferation of nuclear, chemical and biological weapons. In addition, the resolution also requires measure with respect to means of delivery of such weapons. There are complementarities between UNSCR 1540 obligations and the internationally legally binding instruments that make up the global non-proliferation architecture. However, when it comes to means of delivery then there is a lack of such legally binding non-proliferation instruments. This has implications regarding the implementation of UNSCR 1540 measures, as national non-proliferation legislative framework is often based on these international legal instruments.

However, a small number of states refer to national space legislation as part of their legal framework through which they aim to meet their UNSCR 1540 national obligations when it comes to means of delivery. This article aims to explore the nature and context of this, and the possible complementarities with a range of both international non-binding mechanisms and binding legal instruments that come from outside of the traditional non-proliferation canon. This is considered in the context of the developing importance of non-state commercial actors in space activities, and the associated potential emerging proliferation risks, particularly in regard to space-based services, which may inadvertently facilitate the use of nuclear, chemical and biological weapons, and their means of delivery.

Resolution 1540(2004)

Resolution 1540 (2004) was adopted by the Security Council on 28 April 2004 under Chapter VII of the UN Charter in respect to the maintenance of international peace and security. The obligations as set out in the resolution are legally binding on UN member states.

UNSCR 1540 came about as a result, along with the Global Partnership Against the Spread of Weapons and Materials of Mass Destruction (Global Partnership), and the Proliferation Security Initiative (PSI), of the adoption of the so called “Kananaskis Principles” by the G8 (G7 plus Russia) in 2002. The six Kananaskis Principles are intended to prevent terrorists, or those that harbour them, from acquiring or developing nuclear, chemical, radiological and biological weapons; missiles; and related materials, equipment and technology. The resolution itself draws upon the first five of the six principles as given below (Tobey, 2018):

1. Promote the adoption, universalization, full implementation and, where necessary, strengthening of multilateral treaties and other international instruments whose aim is to prevent the proliferation or illicit acquisition of such items; strengthen the institutions designed to implement these instruments.

2. Develop and maintain appropriate effective measures to account for and secure such items in production, use, storage and domestic and international transport; provide assistance to states lacking sufficient resources to account for and secure these items.

3. Develop and maintain appropriate effective physical protection measures applied to facilities which house such items, including defence in depth; provide assistance to states lacking sufficient resources to protect their facilities.

4. Develop and maintain effective border controls, law enforcement efforts and international cooperation to detect, deter and interdict in cases of illicit trafficking in such items, for example through installation of detection systems, training of customs and law enforcement personnel and cooperation in tracking these items; provide assistance to states lacking sufficient expertise or resources to strengthen their capacity to detect, deter and

interdict in cases of illicit trafficking in such items, for example through installation of detection systems, training of customs and law enforcement personnel and cooperation in tracking these items; provide assistance to states lacking sufficient expertise or resources to strengthen their capacity to detect, deter and interdict in cases of illicit trafficking in these items.

5. Develop, review and maintain effective national export and transshipment controls over items on multilateral export control lists, as well as items that are not identified on such lists but which may nevertheless contribute to the development, production or use of nuclear, chemical and biological weapons and missiles, with particular consideration of end-user, catch-all and brokering aspects; provide assistance to states lacking the legal and regulatory infrastructure, implementation experience and/or resources to develop their export and transshipment control systems in this regard.

6. Adopt and strengthen efforts to manage and dispose of stocks of fissile materials designated as no longer required for defence purposes, eliminate all chemical weapons, and minimize holdings of dangerous biological pathogens and toxins, based on the recognition that the threat of terrorist acquisition is reduced as the overall quantity of such items is reduced.

The text of the resolution adopts these principles via binding operative paragraphs that require the national implementation of the required measures. The resolution does reflect principle 1, in recognizing that "...most States have undertaken binding legal obligations under treaties to which they are parties, or have made other commitments aimed at preventing the proliferation of nuclear, chemical or biological weapons",[1] and these include[2]:

- Nuclear Non-proliferation Treaty (NPT)
- Nuclear Weapons Free Zones/Protocols
- International Convention for the Suppression of Acts of Nuclear Terrorism
- Convention on Physical Protection of Nuclear Material (CPPNM)
- 2005 Amendment to the CPPNM
- Comprehensive Nuclear-Test-Ban Treaty (CTBT)
- Chemical Weapons Convention (CWC)
- Biological and Toxin Weapons Convention (BTWC)
- Geneva Protocol of 1925
- 1997 International Convention for the Suppression of Terrorist Bombings
- 1999 International Convention for the Suppression of the Financing of Terrorism
- 2005 Protocol to the Convention for the Suppression of Unlawful Acts Against the Safety of Maritime Navigation
- 2005 Protocol to the Protocol for the Suppression of Unlawful Acts Against the Safety of Fixed Platforms Located on the Continental Shelf
- 2010 Convention on the Suppression of Unlawful Acts Relating to International Civil Aviation
- Other legally binding regional instruments, and other international relevant arrangements.

In practice, what UNSCR 1540 requires is that states take the non-proliferation obligations under the above complementary legally binding international instruments, to which the states have themselves agreed to, and ensure that through national legal frameworks, and as appropriate, that those obligations are also placed on non-state actors under their jurisdiction, as illustrated in figure 1.



Figure 1 UNSCR 1540 & Non-proliferation Architecture Complementarity

What is missing from the above list of binding legal instruments is any specific instrument regarding the "means of delivery". Despite a number of attempts and much discussion over the years the current international non-proliferation framework is still lacking such a binding instrument when it comes to the issue of means of delivery for nuclear, chemical and biological weapons.[3] This is an issue that will be explored further in the next section.

1. S/RES/1540 (2004) Security Council Resolution 1540 (2004) Adopted by the Security Council at its 4956th meeting on 28 April 2004 (page 2)
2. This list is reflected in the 1540 Committee's National Implementation Matrices. And as noted on the Committee's website - Committee matrix template has been subject to periodic revisions since its introduction in 2005. Revised templates were approved by the Committee in 2013 and 2017. https://www.un.org/en/sc/1540/national-implementation/1540-matrices
3. The Chemical Weapons Convention in its definition of "Chemical Weapons" under Article II does reference munitions and devices, that release chemical agents, and also there is further mention equipment for the employment of such munitions and devices, which an be taken to refer to "means of delivery".

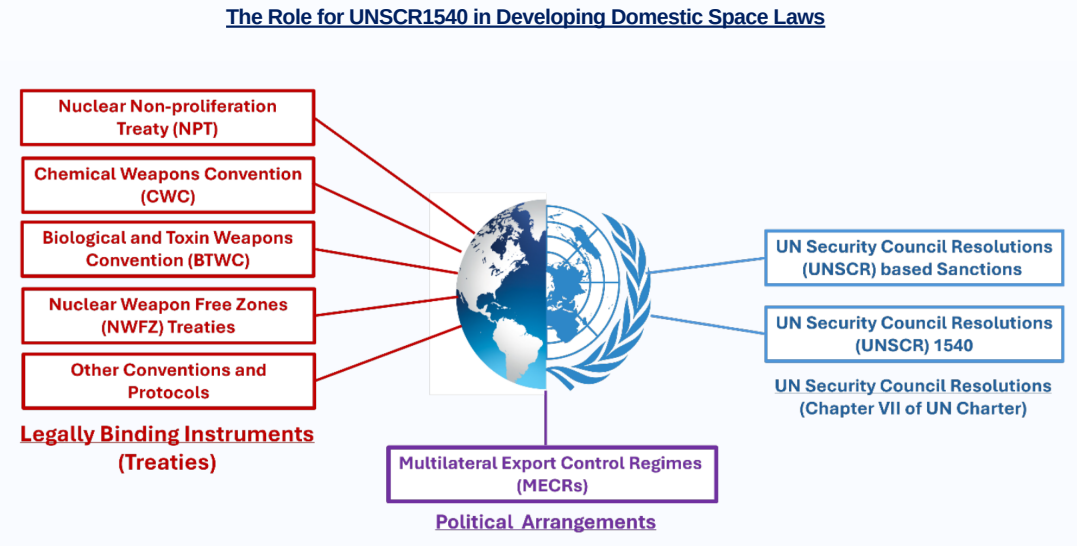

Figure 2 Non-proliferation Framework

However, there are a number of non-legal politically binding arrangements that also form an important part of the non-prolifer-ation framework and are relevant to the implementation of UNSCR 1540. In particular, the multilateral export control regimes (MECRs), figure 2, are complementary to the resolution. The five MECRS are:

- Australia Group (AG), which is concerned with chemical and biological relevant dual-use items;
- Nuclear Suppliers Group (NSG), which is focus on the nuclear fuel cycle and nuclear related dual-use items;
- Zangger Committee (ZC), which is specifically focused on harmonising the interpretation of Article III, Paragraph 2 of the NPT, in respect to nuclear exports[4];
- Wassenaar Arrangement (WA), which is focused primarily on conventional military and related dual-use items, but never-theless there is a cross over with WMD; and
- Missile Technology Control Regime (MTCR), which is specifically focused on WMD means of delivery

These regimes play an important role in regard to the implementation of the resolution in reference to what is referred to as "related materials", as these sensitive dual-use items are normally identified and listed by the relevant MECR. Whilst the focus of MECRs are exports, the lists they produced are often used in wider policy and regulatory areas; this is because the basis for many of the regimes are the multilateral treaties, such as the NPT, CWC and the BTWC.
The key regime, when it comes to the issue of the means of delivery, is the MTCR which is focused on the non-proliferation of unmanned WMD delivery vehicles, such as: ballistic missiles (BM), cruise missiles (CMs), unmanned aerial vehicles (UAVs), space launch vehicles (SLVs) and sounding rockets (SRs), and related dual-use technology. However, the Wassenaar Arrange-ment also controls much of the same systems and related items and technology, but within the context of conventional weapons rather than specifically WMD.

UNSCR 1540 and "means of delivery"

On the specific question as to what UNSCR 1540 means when it refers to "means of delivery" then the resolution provides a specific definition as a footnote, along with the definitions for non-state actors, and related materials see figure 3. For the resolu-tion the "means of delivery" are "missiles, rockets and other unmanned systems capable of delivering nuclear, chemical, or biological weapons, that are specially designed for such use". The resolution is very specific in that its concern is only with unmanned means of delivery and not with either manned or more rudimentary improvised devices. (Whiteside, 2007)

4. Zangger Committee website https://zanggercommitt.org



Non-State actor : individual or entity, not acting under the lawfulauthority of any State in conducting activities which come within the scope of this resolution.



Related materials : materials, equipment and technology covered byrelevant multilateral treaties and arrangements, or included on nationalcontrol lists, which could be used for the design, development, productionor use of nuclear, chemical and biological weapons and their means ofdelivery.



Means of delivery : missiles, rockets and other unmanned systemscapable of delivering nuclear, chemical, or biological weapons, that are specially designed for such use.

Figure 3 UNSCR 1540 (2004) Definitions

When it comes to meeting the obligations of UNSCR 1540 regarding means of delivery, as with other areas, the resolution is not prescriptive as to what states need to do and is therefore reliant on the measures that states have adopted in respect to other international obligations to which they have committed, either via international binding legal instruments or politically binding mechanisms. Specifically in the case of means of delivery, in the absence of legal binding instruments, there is a lot more empha-

WMD Means of Delivery Non-proliferation Mechanisms

As has been noted, with there being no specific internationally legally binding instruments when it comes to WMD means of delivery, around which states can build a national regulatory framework, the focus in this area is on political arrangements such as the MTCR and the Hague Code of Conduct. However, there are some issues with this, for example the MTCR is focused on the export control of sensitive items and technology, whilst the Hague Code of Conduct against the Ballistic Missile Proliferation (HCoC) is a transparency and confidence measure and is only focused on one type of delivery system[5].

Before discussing the MTCR and HCoC in detail, it is worth touching on why there is no multilateral international instrument, and the efforts that have been made in trying to establish such an instrument. Whilst instruments such as the Arms Trade Treaty (ATT) does include missiles within its scope, it doesn't specifically address the issue of missiles, rather treats them the same as other weapon systems. The UN General Assembly established a Panel of Governmental Experts to consider the "issue of missiles in all its aspects", which met a total of three times, July 2001 – July 2002, in 2004, and finally in June 2007 – June 2008. The report from the final panel contained a number of conclusions and recommendations, among these were[6]:

> "… that it would be useful to try to develop understandings on how to categorize and describe missiles and how their different aspects (strategic, political, technical, economic and commercial) interrelate and their impact on global and regional security scenarios."

> And, "… that that step-by-step approach could also include, among other things, efforts to further refine existing national measures of control over transfer and export of missiles and related items, including technology; reporting by States, as appropriate, of missile-related information under the various reporting mechanisms of the United Nations; and efforts to enhance global and regional security, including peaceful settlement of disputes and the promotion of peaceful uses of outer space."

The 2008 report noted concern with the "transfer to and use of certain types of missile and missile technology by State or non-State actors…". However, the work of the three panels was arguably hampered by a major problem, and as the report notes "… the complexity of the issue of missiles, which, besides the various different technical aspects, also had strategic, political, economic and commercial implications, had been complicated by divergent perceptions of the nature of the concern." And that "…these aspects made it difficult for the Panel to reach consensus on many missile-related issues, both fundamental and specific".

5. The HCoC originally was the International Code of Conduct against Ballistic Missile Proliferation, however, the name of Hague Code of Conduct has become the default name through common usage.
6. Report of the Panel of Governmental Experts 28 July 2008 A/63/176 (Experts, 2008)

Whilst attempts at an international binding legal instrument remain to be solved, there are the two key politically binding instruments in the MTCR and HCoC.

Missile Technology Control Regime (MTCR)

The MTCR was established in 1987 because of the growing concern, particularly by G7 countries, about the proliferation of ballistic missiles; this concern was largely driven by the use of such missiles by both sides in the 1980 – 88 Iran – Iraq War. This conflict also saw the largest use of WMD, in the form of chemical weapons particularly by Iraq since the first world war, and included for the first time the use of nerve agents (Tabun and Sarin)[7][8].Whilst the original focus of the MTCR was nuclear weapon delivery systems, in 1992 the MTCR expanded its focus to cover both chemical and biological weapons, this was reflected in the regimes growing interest in UAVs in addition to ballistic and cruise missiles. Currently there are 35 participating states in the MTCR, with an additional five non-participating states as self-declared unilateral adherents.

The MTCR is principally concerned with ballistic missiles, space launch vehicles, sounding rockets and unmanned air vehicle systems, including cruise missiles systems, target and reconnaissance drones. Of particular concern to the regime is the transfer of those systems capable of a range exceeding 300km; with special attention paid to those systems, known as Category-I, which also have a payload capability of 500kg or greater. (MTCR, 2026) The MTCR also controls major-subsystems, components, and a range of related equipment, material and technologies. When it comes to MTCR Category-I systems and subsystems, there is normally a "strong presumption of denial" by MTCR participating states when it comes to their transfer, as well as a total prohibition on the transfer of production facilities for those systems and sub-systems.

When it comes to the issue of space activities the MTCR Guidelines states that the "Regime is not designed to impede national space programs or international cooperation in such programs as long as such programs could not contribute to delivery systems for weapons of mass destruction." In addition, the guidelines also note that in regard to "SLV equipment and technology transfers, however, since the technology used in an SLV is virtually identical to that used in a ballistic missile, which poses genuine potential for missile proliferation". (MTCR, 2026)

The MTCR Guidelines are political, and not legally binding on the regime's participating states, as it is not a treaty or similar international legal instrument. States, may through their domestic legislation, choose to legally bind themselves to the Guidelines. Whilst this is most often done through national export control legislation and policies, however, as will be seen later, other relevant national legislation may also refer to the MTCR Guidelines.

Hague Code of Conduct against Ballistic Missile Proliferation (HCoC)

Whilst the establishment in 2002, of what was known at the time as the International Code of Conduct against Ballistic Missile Proliferation, the Hague Code of Conduct (HCoC) was closely intertwined with the MTCR, but its actual origins go further back to 1987 and in particular the conclusion of the Intermediate-Range Nuclear Forces (INF) Treaty between the United States and the Soviet Union. One idea that came out of the INF was the idea of globalising the Treaty and introducing arms control measures on weapons systems capable of a range between 500 and 5,000kms. Although, nothing eventually came of this particular idea it did lay a foundation to build upon in respect of an arms control approach to ballistic missile proliferation that would be complementary to the export focused MTCR. (Van Diepen, 2022)

The HCoC, like the MTCR, is a non-binding political instrument and has currently 145 subscribing states, initially there were 93 states that signed up for the code. The code has adopted the following general measures:

1. *To ratify, accede to or otherwise abide by:*
   - *the Treaty on Principles Governing the Activities of States in the Exploration and Use of Outer Space, including the Moon and Other Celestial Bodies (1967)*,
   - *the Convention on International Liability for Damage Caused by Space Objects (1972), and*
   - *the Convention on Registration of Objects Launched into Outer Space (1975);*

2. *To curb and prevent the proliferation of Ballistic Missiles capable of delivering weapons of mass destruction, both at a global and regional level, through multilateral, bilateral and national endeavours;*

3. *To exercise maximum possible restraint in the development, testing and deployment of Ballistic Missiles capable of delivering weapons of mass destruction, including, where possible, to reduce national holdings of such missiles, in the interest of global and regional peace and security;*

7. The Iranian casualties during the war represents the third highest numbers of casualties, after the Holocaust and the chemical warfare of World-War I (Ali, 2001)
8. The use of chemical weapons during the Iran – Iraq War, lead to the establishment of the Australia Group (AG) to address the issue of the export of dual-use chemicals and equipment used to develop and produce chemical agents and not directly listed by the CWC. (Ali, 2001)

4. *To exercise the necessary vigilance in the consideration of assistance to Space Launch Vehicle programmes in any other country so as to prevent contributing to delivery systems for weapons of mass destruction, considering that such programmes may be used to conceal Ballistic Missile programmes;*

5. *To contribute to, support or assist any Ballistic Missile programme in countries which might be developing or acquiring weapons of mass destruction in contravention of norms established by, and of those countries' obligations under international dissarmament and nonproliferation treaties.*

The Developing Challenge of "New Space"

Historically the concern with missile proliferation was focused on the state, as reflected in the nature of both the MTCR and HCoC, as systems of concern were primarily developed, manufactured and used exclusively in the context of state actors such as militaries, space agencies, or aerospace companies working under government contract. However, the recent move towards a greater entrepreneurial private sector commercially focused "ecosystem" in space exploration and commercialisation, (sometimes referred to as "New Space" or "Space 2.0"), has shifted the paradigm when it comes to space activities. As a result, a growing number of non-state actors either have access to, or are holders of, sensitive ballistic missile relevant technologies or are potential providers of sensitive space-based capabilities or downstream services. This shift can be characterised as presenting two possible areas of potential concern, these are:

1. Direct support to the development and production of delivery systems such as ballistic missiles, as well as the emerging threat from hypersonic glide vehicles (HGVs) and potentially hypersonic cruise missiles (HCMs).
2. Supporting the operational deployment (use) of WMD delivery systems be they ballistic, cruise missiles or long-range unmanned aerial vehicles.

Direct Support

The potential risks around the use of space launch vehicles (SLV) as a WMD delivery system is well established, due to the commonality in SLV and BM technology, and the shared development history of many systems. (Maitre, 2022) For example, intercontinental ballistic missiles have often been used as SLVs, SLV stages, or as the basis for the development of families of SLVs. (Mistry, 2012) Recent developments in commercial space activities include the development of private space launch systems and capabilities which is potentially an area of increasing proliferation concern. This is because the commercial space sector is increasingly developing and using technologies that have traditionally been under direct state control as highly sensitive technology including MTCR Category-I technologies.

An example of the potential risk around advanced dual-use commercial space technology is that of the potential development of so called high-energy re-usable vehicles, as this technology is closely related to that required for long-range ballistic missile re-entry vehicles (RVs). Potentially, such technology is also useful for the development of advanced non-ballistic systems such as hypersonic glide vehicles (HGVs) and hypersonic cruise missiles (HCMs).

The development of RVs, HGVs and HGMs can potentially use or benefit from a number of commercial space related technologies, these can include:

- Sub-orbital Sounding Rockets (SRs) – the use of SRs to develop strategic weapons has a long-history. For example, the UK's Black Knight was a sub-orbital vehicle designed to test RV technology for the UK's Blue Streak Intermediate-Range Ballistic Missile (IBRM) programme, (specifically the RV radar signature during re-entry), during the late 1950s and early 1960s. Since then, SRs have been used to test hypersonic vehicle technology such as scram jet engine combustion, engine intake design and other aspects of HCM technology, and as the boost stage for HGV testing. A more recent example is the US - New Zealand commercial space launch company Rocket Labs which offers a programme of launches called HASTE (Hypersonic Accelerator Suborbital Test Electron) using a sub-orbital variant of their Electron SLV, which is an MTCR Category-I vehicle. Whilst this programme is primarily aimed at supporting US Government programmes in the development of RV, and hypersonic vehicles, the HASTE launches are also marketed for testing commercial high-energy re-usable vehicle prototypes and subsystems.[9]

- Commercial Re-use/Re-entry Vehicles – the development of commercial re-usable space-launch or space-return vehicles is a key goal in reducing the cost of space access, see figure 4. This development includes the commercial interest in the development of both so called "lower-energy systems", such as re-usable rocket 1st stages, or sub-orbital vehicles, and of "high-energy" reusable space vehicles, these are systems designed to return from a period in orbit. As noted, the technologies required for such high-energy re-usable systems are very similar to that required for HGVs and potentially HCMs, as well as BM RVs, and is controlled under MTCR Category-I.

9. https://rocketlabcorp.com/launch/haste/

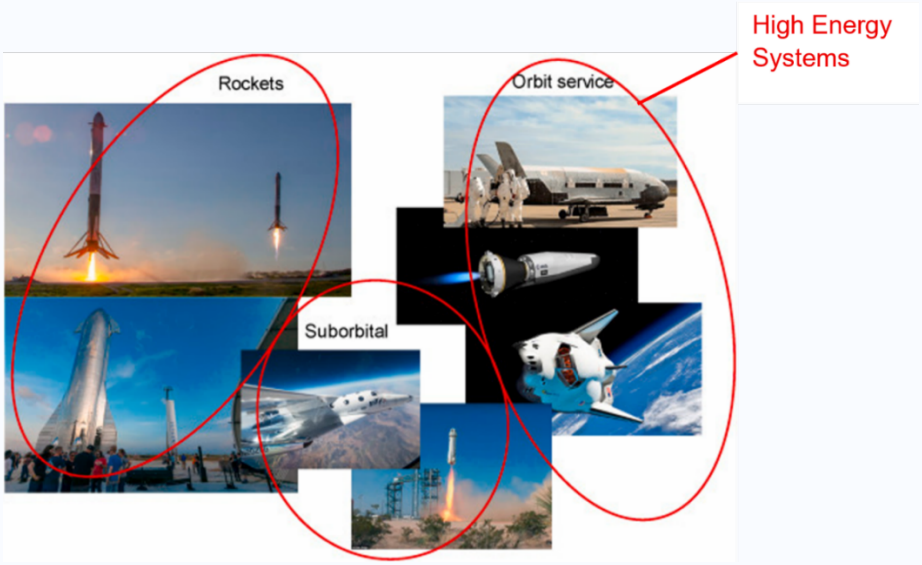

Figure 4 Examples of new generation re-usable rockets and space systems

(Image Credit: (Baiocco, 2021))

One potential commercial space application that will require high-energy re-entry systems is that of on orbit services, such as space-based manufacturing or logistics that require the return to earth of items manufactured or deployed in orbit, several commercial companies are developing concepts around this return-to-earth commercial requirements.

The actual difference in the technology between a commercial RV, military RV and HGV is minimal. For example, the fundamental difference between a ballistic missile RV, a space return system such as a space capsule, and an HGV is the amount of aerodynamic lift generated as opposed to the drag also generated, (termed as the lift/drag ratio (L/D)). The lift/drag ratio is mostly a function of the aerodynamic shape of the vehicle body, see figure 5.

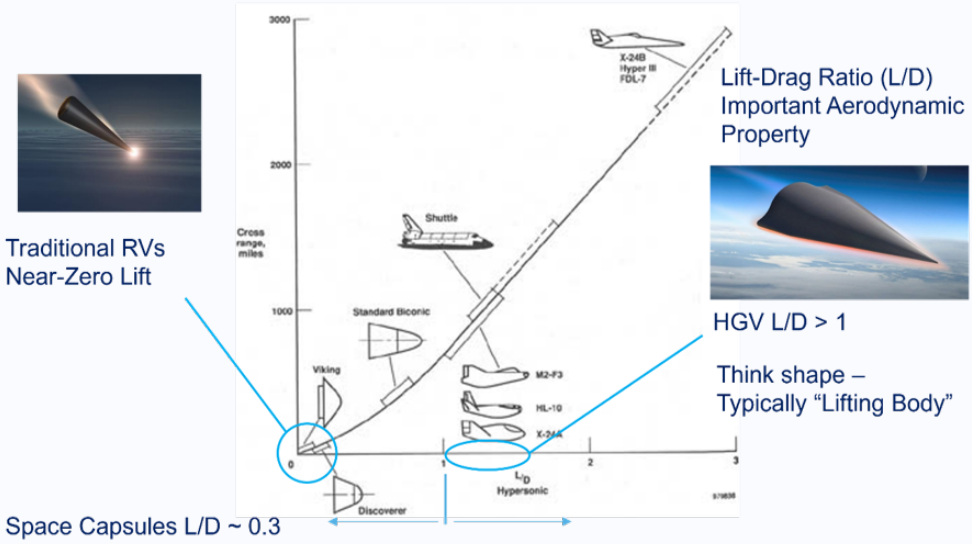

Figure 5 Simple Difference in RV and HGV

Supporting the Use of WMD Delivery Systems

The traditional linkage between SLV and BM technologies is reflected in the MTCR and the HCoC. The MTCR, in being focused on the risks around missile relevant technology transfers, is less focused on possible proliferation relevant activities, particularly when it comes to the use of delivery systems. Whilst the HCoC, coming from a more arms control background as a transparency and confidence building measure, is more focused on activities, and is aimed at reducing the risk of space launch programmes being use as a cover for ballistic missile development. However, the HCoC has traditionally been focused on the state, rather than the potential risks around non-state commercial space actors and is primarily only focused on launch and not proliferation relevant activities undertaken by space-based systems.

An area of potential concern is around the use of a range of space-based systems to support the operational deployment / use of a range of strategic systems, not just BMs, but also cruise missiles (including hypersonic

systems) and long-range UAVs. The recent RAND research report – Exploring Duality in Space: How Select Countries Approach Dual-Use Space Systems, provides a useful summary of the potential uses that a range of space-based dual-use systems, see figure 6.

Dual-Use Capability	Description	Possible Civil Application(s)	Possible Military Application(s)
Longstanding dual-use integration			
Satellite communications (**SATCOM**)	Use of orbiting satellites to enable data, voice, and video transmission over long distances	Internet and phone service, television and radio broadcasting, connectivity services	Command and control, surveillance, operational flexibility
Positioning, navigation, and timing (**PNT**)	Global navigation satellite systems (GNSS) providing timing and positional data for navigation and communication	Transportation safety, autonomous vehicle navigation, positional awareness, economic infrastructure	National security operations, resilient communication networks
Remote sensing for Earth observation (**EO**) and space-based intelligence, surveillance, and reconnaissance (**ISR**)	Satellite systems that gather data about Earth's surface and monitor activities globally	Precision agriculture, disaster response, climate monitoring, water resource management	Intelligence-gathering, situational awareness, threat detection, imaging, operational and tactical military support

Figure 6 Extract from Table 1.1 (page 5) of the RAND Research Report

Exploring Duality in Space

Figure 7 illustrates the potential strategic application of the various highlighted space-based systems mentioned in the RAND report. Such commercial systems have been used to augment military capabilities, and to fill in for either a lack of a specific capability or to supplement existing systems by ensuring continued geographical or temporal coverage. However, this situation could also apply to violent non-state actors, such as terrorist or insurgent groups, who can use space-based services to support the planning and targeting, even if deploying more rudimentary weapon delivery means.

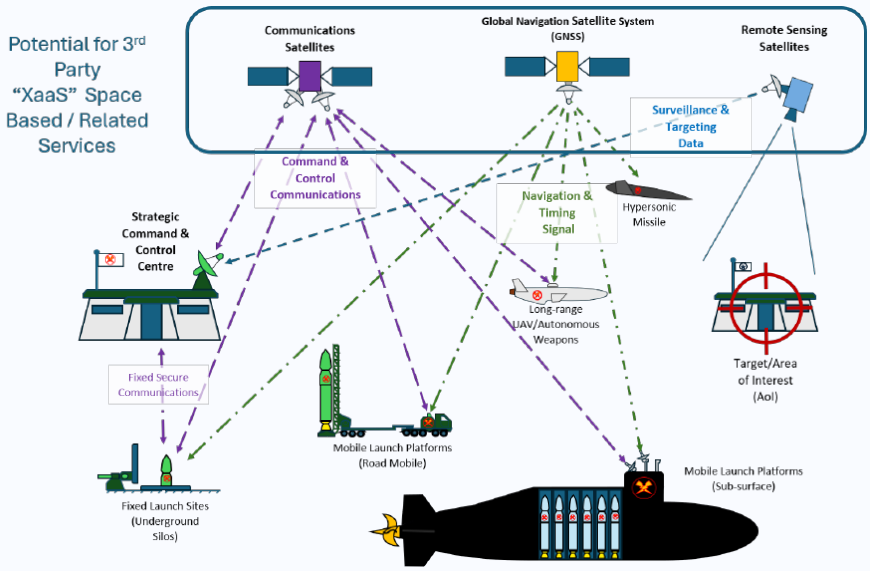

Figure 7 Strategic Weapons Systems

As has been noted the traditional approach to non-proliferation has focused on the transfer of either physical items or technology, particularly for the development and production of systems. There has been less focus on the provision of a capability or service especially where there is no transfer of hardware, software or controlled technology required to access that capability.

Commercial space sector activities are often referred to as either "upstream" - the development and provision of space technology; or "downstream" which is focused on providing a service or the utilization of space data, see figure 8. It is the downstream provision of a capability or service that traditionally has not been considered in the context of non-proliferation. However, over recent years there has been some discussion within the multilateral regimes around the emerging concept of "X as a Service" (XaaS). This has been focused on cloud-based services, such as data storage and cloud computing, but the concept can be applied to downstream space-based services such as SATCOM, PNT and EO/ISR.

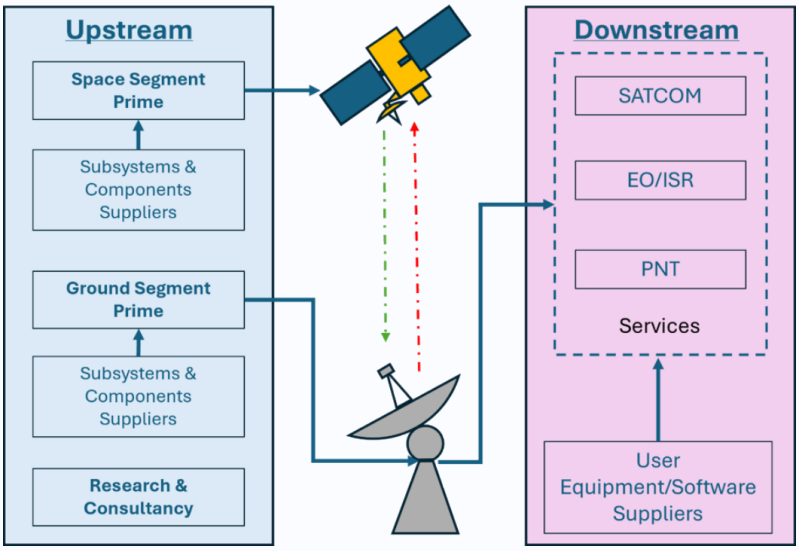

Figure 8 Commercial Space Sector "Upstream" and "Downstream"

Dual-Use "X as a Service"

As has been mentioned the potential application of dual-use space systems can provide vitally important support to the operational deployment of strategic weapon systems, as illustrated in figure 6. However, the concept of "dual-use", even in the absence of an agreed definition, is well-established regarding traditional strategic export controls with its focus on technology. However, it is arguably less well understood when applied to the provision of services such as "XaaS".

Everything (X) as a Service, has developed as a business model over the last few years, and is based on the concept of the user being less dependent on the acquisition of specific hardware or infrastructure, but is able to have access to the required capability. Typically, XaaS business models are cloud or online based and focused on data. A good example of XaaS in the space context is that of *Commercial Satellite Imagery (CSI)* as a source of EO/ISR data.

One example of the potential issue around CSI is highlighted by a number of media reports regarding US CSI providers, such as Planet Labs and Vantor (formally Maxar), who supposedly at the request of the US Government, are now operating what is termed a "controlled accesses model" or "enhanced access controls" regarding the release of commercial satellite imagery of Iran and the wider region.[10] Whilst the available, relatively low-resolution commercial imagery, may not be that useful for precision targeting purposes, never-the-less such imagery is useful for things like Battle Damage Assessment (BDA) which would allow for the adjustment and refine any ongoing missile/drone campaign to ensure effective and efficient use of weapon stocks, and for monitoring commercial maritime activities.

"CIS as a service", allows the end-user to specify the area of interest (AoI) to be selected, as well as a number of other possible collection requirements, all done via the internet, this is referred to as tasking, see figure 9.

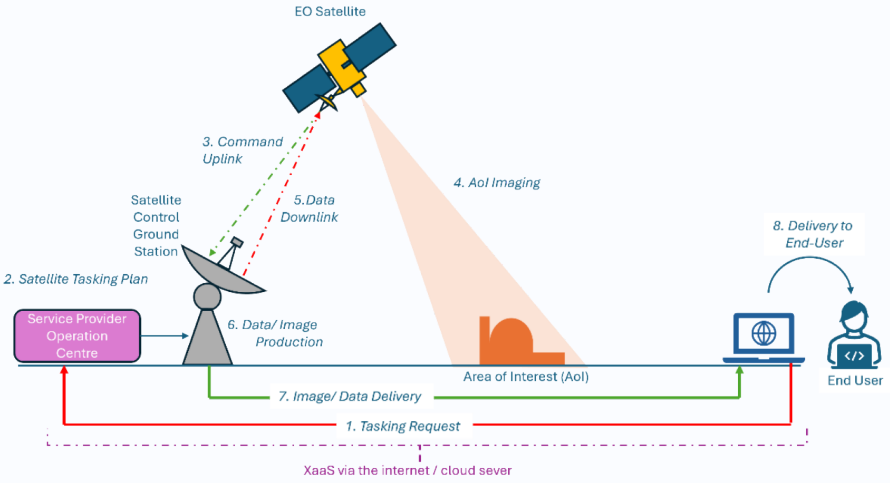

10. An example of the reporting is the "Pentagon Enforces Commercial Satellite Blackout Over Middle East War Zone," Satnews, 20 May 2026. It should be noted that there is no confirmation that the U.S. Government triggered the "Interruption of Normal Commercial Operations", (so called "shutter controls"), provided for under U.S. commercial operators' licensing requirements.

Figure 9 CSI Tasking of Imagery Collection

(based on Airbus tasking process for Pléiades Neo via One Altas)

One potential avenue regarding the control of satellite imagery intended for the use of means of delivery is whether or not a state would regard such imagery data, if specifically tasked and collected for the purpose of missile guidance or navigation, as controlled under national export control legislation. Here the potential for control is based on the "data" being seen as "technical data" within the MTCR definition of "technology". However, this does potentially raise an interesting legal question around the definition of an export, (more correctly an intangible transfer of technology) from a satellite in orbit if the imagery data is downloaded to a ground station in a 3rd country or direct to the customer, see figure 10. Alternatively, such provision of a space-based service, when connected to use of a WMD means of delivery, may be seen as a provision of either a technical service or provision of "technical assistance" which is the other component of the MTCR definition of controlled "technology". (MTCR, 2026) However, basing possible control within, for example, national space regulatory framework would arguably simplify matters, both for the regulator and the service provider, as it would not be dependent on export controls, which as noted, are more focused on traditional transfer of items and not on activities.

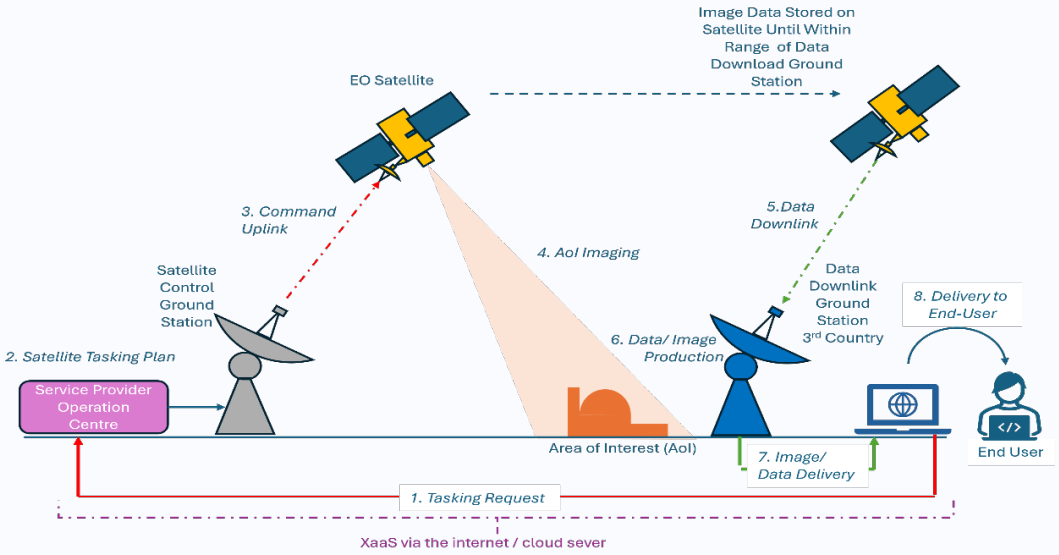

Figure 10 CSI Tasking of Imagery Collection with Image Data in 3rd County

UAVs and Space-Based XaaS

While there is no suggestion that violent non-state actors (VNSAs) would have access to the types of sophisticated strategic weapons systems as presented in figure 7, However, the growth in the use of unmanned aerial vehicles (UAVs) has accelerated over last few years, including their use by VNSAs, and has become an area of concern with the potential for the use of UAVs in the delivery of WMD payloads. (UNOCT, 2024) Whilst the majority of the use of UAVs by VNSA have been at relatively short-range, the potential use of longer-range UAVs cannot be ignored. A key limiting factor regarding the use of longer-range UAV systems is the issue of BVLOS (beyond visual line of sight). The BVLOS issue is due to the curvature of the earth, and other factors, where there is the loss of the direct line of sight between a ground-based control terminal and the UAV, resulting in the loss of the command-and-control datalink. Beyond typically 150km there is a need to have some form of data relay to allow for command-and-control datalink. This is often achieved using SATCOM satellite communication data networks, that typically use non-geostationary satellite orbit (NGSO) constellations of data satellites such as Iridium, OneWeb and Starlink. Increasingly, commercial companies are now offering command-and-control datalink terminals specially designed for UAV BVLOS operations. So, in addition to the potential concern with CSI, the growing potential for space-based data SATCOM services to UAV operations, present a possible risk around the use of long-range UAV means of delivery of WMD payloads.

International Space Law & National Space Regulation

Other than the MTCR and HCoC, and despite a number of attempts, there remains no real multilateral missile focused mechanism, especially not one that takes into consideration the wider proliferation risks around commercial space activities. However, an alternative approach may be to consider the frameworks under which space governance is being developed more generally.

The recent development of commercial space activities has resulted in the growth of national space legislation aimed at regulating commercial space activities. The UN Office of Outer Space Affairs (UNOOSA) maintains a list of

national space laws. The report of April 2024 from the Legal Subcommittee of the *UN Committee on the Peaceful Uses of Outer Space*, listed some 46 states that have implemented some form of space legislation.

Mostly, national space legislation is focused on meeting obligations under the *1967 Outer Space Treaty (OST)*; however, of particular potential relevance to the issue of missile proliferation are OST Articles III and VI:[11]

> Article-III – States Parties to the Treaty **shall carry on activities in the exploration and use of outer space**, including the Moon and other celestial bodies, **in accordance with international law, including the Charter of the United Nations, in the interest of maintaining international peace and security** and promoting international co-operation and understanding.

> Article -VI – States Parties to the Treaty shall bear international responsibility for **national activities in outer space**, including the Moon and other celestial bodies, whether such activities are carried on by governmental agencies or by **non-governmental entities**, and for assuring that national activities are carried out in conformity with the provisions set forth in the present Treaty. **The activities of non-governmental entities in outer space**, including the Moon and other celestial bodies, **shall require authorization and continuing supervision by the appropriate State Party to the Treaty.** When activities are carried on in outer space, including the Moon and other celestial bodies, by an international organization, responsibility for compliance with this Treaty shall be borne both by the international organization and by the States Parties to the Treaty participating in such organization.

Based on the OST Articles III and VI, then there are legally binding obligations placed on states to ensure that any commercial space activity is consistent with international law, including the UN Charter, (such as Security Council resolutions under Chapter VII of the Charter, i.e. UNSCR 1540) in the interests of maintaining international peace and security.

Examples of where national space regulation has recognised such obligations in the context of missile proliferation, includes:

- South Africa's *Space Affairs Act No. 84/1993*, which refers to the "*controlling and restricting the development, transfer, acquisition and disposal of dual-purpose technologies, in terms of international conventions, treaties and agreements*", noting that "*dual-purpose*", as defined in the Act refers to space technologies which can contribute to the proliferation of weapons of mass destruction.
- New Zealand's *Outer Space and High-altitude Activities Act 2017*, which in section 84 specifically refers to the requirement to notify the government of the intention to develop or acquire "missile technology" (in reference to the MTCR).
- Australia's *Space (Launch and Returns) Act 2018*, specifically prohibits launch activities relating to nuclear weapons or weapons of mass destruction.
- United Kingdom's *Space Industry Act 2018*, which includes reference to granting a licence only in line with international obligations, and that the Civil Aviation Authority (CAA), as the regulatory authority, notes that such relevant international obligations include the HCoC.

In particular, much of the relevant domestic space regulations are focused on space launch activities. These types of regulation would clearly be the most appropriate mechanism regarding the implementation of the provision of the HCoC.

When it comes to the specific issue of CSI, then whilst there is no international legally binding instrument, there is the UN General Assembly (UNGA) resolution 41/65 *Principles Relating to Remote Sensing of the Earth from Outer Space* which was adopted on 3rd December 1986. Of note is Principle XIV, which states -

> *"In compliance with Article VI of the Treaty on Principles Governing the Activities of States in the Exploration and Use of Outer Space, including the Moon and Other Celestial Bodies,* ***States operating remote sensing satellites shall bear international responsibility for their activities and assure that such activities are conducted in accordance with these principles and the norms of international law, irrespective of whether such activities are carried out by governmental or non-governmental entities*** *or through international organizations to which such States are parties. This principle is without prejudice to the applicability of the norms of international law on State responsibility for remote sensing activities." (UNGA, 1986)*

The development of specific domestic legislation in respect of the regulation of commercial remote sensing is arguably less well developed than for the more general regulation of space activities, with only a few countries having developed specific regulations so far. These are:

- The US 1984 - *Land Remote Sensing Commercialzation Act*[12]
- Canada 2005 - *Remote Sensing Space Systems Act*

11. 1967 Treaty on Principles Governing the Activities of States in the Exploration and Use of Outer Space, including the Moon and Other Celestial Bodies https://www.unoosa.org/oosa/en/ourwork/spacelaw/treaties/introouterspacetreaty.html

12. The 1984 Act was repleaded by the Land Remote Sensing Policy Act of 1992, see U.S. Code: Title 51 — National and Commercial Space Programs, Subtitle VI Earth Observation, Chapter 601 Land Remote Sensing Policy, Subchapter III Licensing of Private Remote Sensing Systems

- Germany 2007 - *Protection against the Security Risk to the Federal Republic of Germany by the Dissemination of High-Grade Earth Remote Sensing Data (Satellite Data Security Act)*
- France 2008 - *French Space Operation Act*
- Japan 2016 - *Act on Ensuring Appropriate Handling of Satellite Remote Sensing Records*
- Finland 2023 - *Government Decree on Limit Values for Satellite Remote Sensing (amending the 2018 SpaceActivities Act)*

A consideration for national law makers is the competing requirements to ensure national security, including through meeting state's international obligations such as UNSCR 1540, and ensuring that domestic industries are not disadvantaged in a global market by regulations that are too restrictive in comparison to other states. However, a point noted by Buamann and Pellander is that "the more technology advances, the more security concerns may arise".[13] This means that in all likelihood that there will be a greater need for the regulation of commercial remote sensing. Especially, as more state's will have commercial operators under their jurisdiction with evermore sophisticated systems offering increasingly advanced capabilities, which potentially will increase the security concerns around who is accessing those capabilities, including in the context of potential use of WMD by violent non-state actors.

The Role for UNSCR1540 in Developing Domestic Space Laws

The 2020 paper by Kolja Brockmann (Brockmann, 2020) discussed the concept of developing further the complementarity between UNSCR 1540, the MTCR and HCoC, with the aim of strengthening missile non-proliferation. Building on this idea, then a further step may be to consider further complementarity with respect to relevant international space law, such as the OST, through the development of domestic space regulatory frameworks. In practice this would mean national space launch regulations should reflect non-proliferation concerns as an integral element of meeting OST Articles III and VI, for example in line with obligations under UNSCR 1540, as well as relevant national obligations under bilateral treaties (such as Technology Safeguards Agreements regarding the transfer of sensitive technologies), and relevant multilateral arrangements such as the MTCR and HCoC.

Going beyond space launch, then commercial space-based activities that carry a proliferation related risk as discussed above, such as SATCOMs, PNT and EO/ISR, that are subject to national space regulations in line with the OST Articles, then ideally those regulations should also take into account non-proliferation issues. Especially, the provision of downstream services that potentially have higher-levels of risk exposure around supporting the operational deployment / use of WMD armed BM, CM and UAVs. For example, the current EU proposal for a Regulation of the European Parliament and of the Council on the Safety, Resilience and Sustainability of Space Activities in the Union, could include provisions that any space activities undertaken by commercial space entities in the EU do not contribute to the operational deployment of the use of WMD means of delivery.[14] This could be linked to the latest Council Decision (CFSP) 2023/124 of 17 January 2023, in support of the Hague Code of Conduct and ballistic missile non-proliferation in the framework of the implementation of the EU Strategy against Proliferation of Weapons of Mass Destruction.[15][16] This arguably would be a clear demonstration of the potential complementarities between developing national space regulatory frameworks and meeting obligations around UNSCR 1540 means of delivery.

Means of Delivery Issue	International Binding-Legal Instruments	International Non-Binding Mechanisms
Support to the Development / Production of Missile Systems		MTCR (i.e. Export Control of Sensitive Items and Technology)
		HCoC (i.e. Transparency and Confidence Measures around Space Launch Activities)
Support Regarding the Deployment / Use of Missile Systems	Outer Space Treaty Articles III & VI, (i.e., Adherence to International Peace and Security Obligations)	
		UNGA Principles Relating to Remote Sensing Principle XIV (i.e. Adherence to International Peace and Security Obligations via OST III and VI)
		MTCR (i.e. Export Control of Provision of Technical Assistance & Transfer of Technical Data (Satellite Imagery Data for Missile Guidance etc))

Table 1 UNSCR 1540 Missile and Potential Space Related Complementarities

13. (Baumann & Pellander, 2024) page 298
14. Proposal for a REGULATION OF THE EUROPEAN PARLIAMENT AND OF THE COUNCIL on the sfety, reslience and suetainablity of pace ectivties in the Union,Document 52025PCO35https://eur-lex.europa.eu/legal-content/EN/TXT/?uri=celex:52025PC0335
15. Council Decision (CFSP) 2023/124, https://eur-lex.europa.eu/egal-content/EN/TXT/PDF/?uri=CELEX:32023D0124.
16. EU Strategy Against Proliferation of Weapons of Mass Destruction, https:/data.conslum.europs.eu/doc/document/ST-15708-2003-INIT/en/pdf

Conclusion

Given the lack of legally binding international instruments in regard means of delivery for nuclear, chemical and biological weapons, around which states can develop domestic legislative frameworks to meeting their UNSCR 1540 obligations in implementing measures, means there is an argument for looking at alternatives, including outside of the traditional international non-proliferation architecture, that would provide additional complementarities with respect to the state's implementation of the resolution's obligations, see figure 10. This is even more important given the growing access that non-state actors potentially have to these proliferation sensitive technologies. Furthermore, there is arguably a case for also considering moving beyond the traditional view of potential concerns as simply about the direct contribution that, for example, space launch vehicle technology can make to the development and production of ballistic missiles. Certain areas of commercial space interest such as the development of high-energy re-useable vehicles can also be seen to include potential concerns around the contribution developing such dual-use technology can make to the development of hypersonic glide vehicles and hypersonic cruise missiles systems.

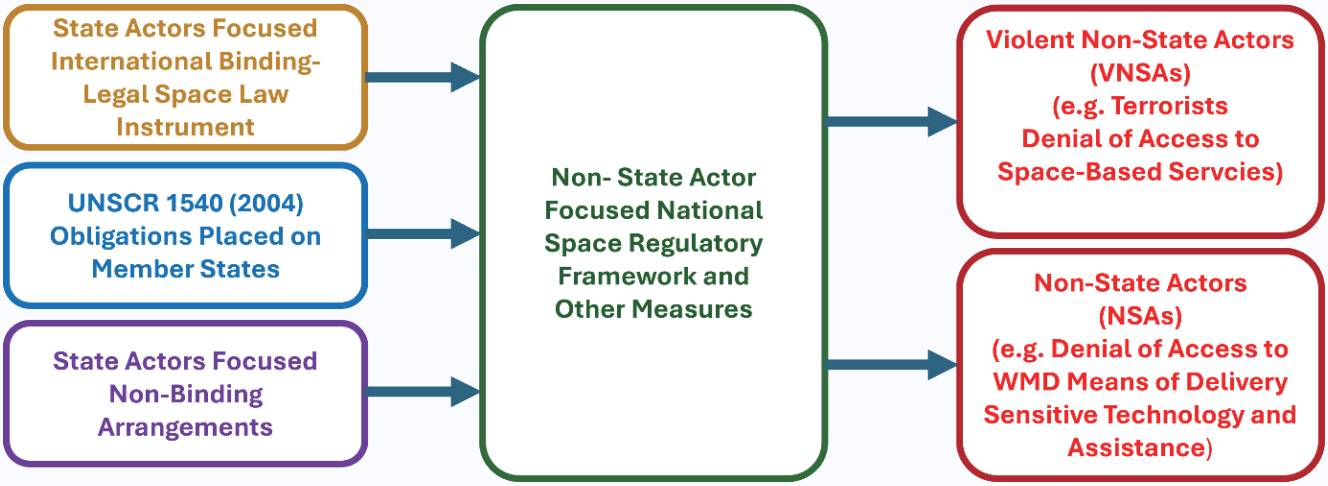

Figure 11 UNSCR 1540 Complementarities

Further consideration is also needed in regard to the possibility of commercial space-based systems and services and their potential contribution to the deployment and use of WMD means of delivery. The security risks in this area are especially of concern when considered in conjunction with business practices that are becoming more focused on virtual delivery via the internet, such as tasking by the customer and direct data delivery to the end-user of sensitive data from satellites, (or via a 3rd country), and so bypassing traditional means of regulation such as export controls.

When developing national space legislative frameworks and domestic regulations regarding non-state commercial space activities, then ideally states should also consider the complementarities of these regulations with their UNSCR 1540 obligations, as these measures could be useful in addressing the means of delivery issue, as noted there are a few states that have already done this.

The Author

Andrew Horton is a visiting senior research fellow at King's College London's Centre for Science and Security Study (CSSS). He has over 20-years of experience with the UK Government in various roles in strategic export controls, non-proliferation and arms control, including UNSCR 1540. He is a former chair/co-chair of the Missile Technology Control Regime (MTCR) and for over ten-years UK head of technical delegation to the MTCR and Nuclear Suppliers Group (NSG). As part of his work in the UK Government, he helped to draft and develop the UK national space regulatory framework working closely with the UK Space Agency and the UK space sector. He holds a degree in Aerospace Engineering, MA in Science & Security, and an MSc in Space Science & Technology.

Bibliography

Ali, J. (2001). Chemical Weapons and the Iran-Iraq War: A Case Study in Noncompliance. The Nonproliferation Review (Spring), 43 - 58.

Baiocco, P. (2021). Overview of reusable space systems with a look to technology aspects. Acta Astronautica,, 189, 10 - 25.

Baumann, I., & Pellander, E. (2024). Regulation of Commercial Earth Observation Systems and Data . In B. I. Smith L. (Ed.), Routledge Handbook of Commercial Space Law (pp. 297 - 314). Routledge .

Brockmann, K. (2020). Controlling Ballistic Missile Proliferation: Assessing Complementarity Between the HCoC, MTCR and UNSCR 1540. HCoC Research Papers No.7.

COPUOS, C. o.-L. (2024). Schematic Overview of National Regulatory Frameworks for Space Activities A/AC.105/C.2/2024/CRP.10 .

Experts, P. o. (2008). A/63/176 Report of Secretary-General - The Issue of Missiles in All Its Aspects . New York: UN General Assembly .

Maitre, E. M.-B. (2022). The Hauge Code of Conduct and Space. HCoC Research Papers (No.10).

Mistry, D. G. (2012). Ballistic Missiles and Space Launch Vehicles in Regional Powers. Astropolitics, 10(2), 126 - 151.

MTCR. (2026). MTCR Guidelines. Retrieved 07 10, 2026, from https://www.mtcr.info/en/mtcr-guidelines

Tobey, W. (2018). A History of United Nations Security Council Resolution 1540. In D. S. Salisbury (Ed.), Preventing the Proliferation of WMDs (pp. 13 - 32). Palgrave Pivot, Cham.

UNGA. (1986). UNGA 41/65 Principles Relating to Remote Sensing of the Earth from Outer Space.

UNOCT. (2024). Global Report on the Acquisition, Weaponization and Deployment of Unmanned Aircraft Systems by Non-State Armed Groups for Terrorism related Purposes. New York : United Nations Offcie of Counter-Terrorism (UNOCT) Global Counter-Terrorism Programme on Autonomous & Remotely Operated Systems (AROS Programme).

Van Diepen, V. H. (2022). Origins and Development of the Hague Code of Conduct. HCoC Papers, No. 11.

Whiteside, T. (2007). UNSCR 1540 and "Means of Delivery". In O. &. Bosch (Ed.), Global Non-Proliferation and Counter-Terrorism : The Impact of UNSCR 1540 (pp. 113 - 121). New York: Bloomsbury Publishing USA,.

ARTICLE 03

Artificial Intelligence as a Tool for Preventing the Proliferation of Weapons of Mass Destruction: A Civilian Perspective

Sai Gu
Deputy Pro-Vice-Chancellor (China) and Professor of Chemical and Process Engineering
University of Warwick



Abstract

United Nations Security Council Resolution 1540 (2004) established a binding framework to prevent non-State actors from acquiring nuclear, chemical, and biological weapons. Two decades on, the rapid advancement of artificial intelligence presents both new proliferation risks and novel opportunities for strengthening this regime. This article examines how AI applications—across machine learning, deep learning, natural language processing, and computer vision—can support civilian-led non-proliferation efforts in five key areas: verification and monitoring, prediction and early warning, export controls and supply chain security, detection of chemical and biological threats, and international cooperation and capacity-building. Drawing on recent developments at the IAEA, OPCW, DARPA, and other international bodies, the article provides a structured analysis linking specific UNSCR 1540 operative paragraphs to concrete AI deployment examples via a systematic mapping table. It argues that AI, subject to appropriate governance frameworks, sustained international cooperation, and

meaningful human oversight, can serve as a valuable complement to traditional non-proliferation mechanisms. The article is framed as a contribution to academic awareness, education, and implementation-support, reflecting the guidelines of the UNSCR 1540 TechHorizon publication project.

Keywords: UNSCR 1540; artificial intelligence; non-proliferation; weapons of mass destruction; dual-use technology; export controls

1. Introduction

United Nations Security Council Resolution 1540, adopted unanimously in 2004 under Chapter VII of the UN Charter, established a binding international framework to prevent non-State actors from acquiring nuclear, chemical, and biological weapons and their means of delivery. The resolution obligates all States to adopt and enforce effective measures against the proliferation of such weapons, including the enactment of domestic legislation criminalising proliferation activities, the establishment of robust export and border controls, and the physical protection of sensitive materials and facilities (United Nations Security Council, 2004). The resolution also established a dedicated committee—the 1540 Committee—supported by a Group of Experts, to monitor implementation and facilitate assistance to States in meeting their obligations. As a resolution adopted under Chapter VII, its provisions are legally binding on all UN Member States, making it a uniquely powerful instrument in the international non-proliferation architecture.

Twenty years after its adoption, UNSCR 1540 remains the cornerstone of the international community's efforts to address the proliferation of WMD to non-State actors. However, the technological landscape has transformed dramatically since 2004. The rapid advancement of artificial intelligence, machine learning, and related digital technologies has created both new proliferation risks and new opportunities for strengthening the non-proliferation regime. AI systems are increasingly capable of analysing vast datasets, detecting anomalous patterns, automating complex monitoring tasks, and generating insights that would be impossible for human analysts alone to produce at scale. The UN Secretary-General has observed that the nuclear non-proliferation regime must evolve to survive the age of AI and other emerging technologies, a sentiment that applies with equal force across the chemical and biological domains.

AI is itself a dual-use technology of profound significance. The same capabilities that enable AI to strengthen verification, monitoring, and export controls also create new pathways for proliferation. AI can accelerate the design of novel chemical agents, optimise biological weapon development pathways, and help malicious actors identify vulnerabilities in control systems. This duality is not unique to AI—many technologies relevant to WMD proliferation have legitimate civilian applications—but the breadth and pace of AI advancement make the challenge particularly acute. As the National Academies of Sciences, Engineering, and Medicine (2025) have observed in their landmark study on AI in the life sciences, the benefits of AI for scientific discovery and public health must be carefully weighed against the potential for misuse, and governance mechanisms must evolve in parallel with the technology itself. The scale and speed of AI advancement have led international research institutes to produce state-of-the-art assessments of its non-proliferation implications. A comprehensive SIPRI compendium, for example, consolidates expert analysis across the military, chemical, biological and nuclear domains, underscoring both the risks AI poses and the opportunities it offers for strengthening the non-proliferation regime (Reinhold et al., 2025). Similarly, Puscas (2024) has examined how large language models and other AI-enabled tools can affect international security, further illustrating the dual-use character of these technologies (Puscas, 2024).

This article examines the potential of AI applications to support the objectives of UNSCR 1540 from a civilian perspective. It focuses on five key areas: verification and monitoring, prediction and early warning, export controls and supply chain security, detection of chemical and biological threats, and international cooperation and capacity-building. Throughout, it emphasises the importance of maintaining human control and oversight over AI systems used in security-related contexts, and it frames the analysis as a contribution to academic awareness, education, and implementation-support rather than as policy prescription.

UNSCR 1540 Operative Paragraph	Key Obligation	AI Application	Real-World Deployment / Example
OP 1: Ban support to non-State actors seeking WMD	Prevent material/technical support to proliferation networks	AI-driven trade & financial screening to detect supply chains aiding non-State actors	ML classifiers auto-flag export license applications for potential diversion to prohibited users
OP 2: Establish domestic controls to prevent proliferation	Account for and secure nuclear, chemical, and biological materials/facilities	AI-enhanced video surveillance & anomaly detection at sensitive facilities; predictive maintenance	IAEA uses AI/ML to auto-analyze safeguards surveillance footage & satellite imagery for undeclared activity
OP 3(a): Account for and secure sensitive materials in production, use, storage, or transport	Chain-of-custody monitoring for WMD-relevant materials	AI-powered supply chain analytics to track materials, detect inventory anomalies, flag diversion risks	Defense Innovation Unit AI-driven CBRN supply chain risk analytics
OP 3(b): Maintain physical protection for sensitive materials and facilities	Perimeter security, access control, intrusion detection	AI-based intrusion detection, behavioral analytics for insider threats, automated surveillance	DARPA SIGMA+ uses AI-driven networked chemical/biological sensors for wide-area threat detection
OP 3(c): Enforce border controls to prevent illicit trafficking	Border security, detect illicit WMD-related trafficking	AI-enabled detection at borders; ML analysis of shipping manifests and customs declarations	OPCW AI Research Challenge: AI models predict novel toxic compounds to improve border detection
OP 3(d): Establish national export/trans-shipment controls on WMD-related items	Export licensing, end-user verification, catch-all controls	NLP & ML to screen export applications, classify dual-use items, assess end-user risk	BIS policy: AI model training may fall under WMD end-use export controls
OP 3(d) cont'd: Control intangible technology transfers	Prevent knowledge/expertise transfer for prohibited purposes	AI monitors scientific publications, patents, and online platforms for proliferation-relevant knowledge transfer	NDU study: AI can monitor dual-use research and open-source data for biological weapons activity
OP 3(d) cont'd: Controls on funds and services for proliferation financing	Financial screening, detect proliferation financing	AI-driven transaction monitoring to detect patterns linked to procurement networks	ML identifies anomalous financial flows tied to WMD procurement
OP 6: 1540 Committee reports on implementation	Monitor national implementation, analyze reporting	AI-powered analysis of national implementation reports to identify gaps and trends	NLP could auto-analyze 1540 implementation data from national reports
OP 7: States assist other States lacking legal/ regulatory infrastructure	Capacity-building, technical assistance, technology transfer	AI-enhanced training platforms, machine translation, automated screening for resource-limited States	OPCW Shanghai Workshop (June 2025): first AI capacity-building programme for States Parties
OP 8(a): Promote universal adherence to non-proliferation treaties	Strengthen the broader treaty regime	AI tools developed for one treaty (e.g., CWC) adaptable to others (e.g., BWC, NPT)	UN General Assembly resolutions supporting human control over AI in security contexts

OP 9: Don't conflict with peaceful use rights under NPT, CWC, or BWC	Balance non-proliferation with legitimate civilian access	Governance frameworks that enable peaceful AI use while preventing misuse	National Academies consensus study: governance should enable beneficial AI in life sciences while mitigating biosecurity risks

Table 1: Structured mapping of selected UNSCR 1540 operative paragraphs to specific AI capabilities and real-world deployment examples

2. Enhancing Verification and Monitoring Capabilities

Verification is the foundation of any effective non-proliferation regime. Without the means to verify compliance, international obligations risk becoming aspirational rather than binding. The challenge of verification has grown significantly in the decades since UNSCR 1540 was adopted: global trade has expanded dramatically, dual-use technologies have become more widely available, and the volume of data that must be analysed to detect proliferation indicators has increased exponentially. AI offers significant potential to strengthen verification and monitoring activities across the nuclear, chemical, and biological domains, enabling human analysts to focus their expertise where it is most needed.

2.1 Nuclear Verification

The International Atomic Energy Agency (IAEA) Department of Safeguards has been actively exploring ways to leverage AI to enhance the effectiveness and efficiency of its safeguards activities. The IAEA's safeguards system is the world's most technically advanced international verification regime, responsible for monitoring nuclear materials and activities across more than 180 States. The sheer scale of this task—encompassing thousands of nuclear facilities, millions of items of nuclear material, and vast quantities of surveillance data—makes it a natural candidate for AI augmentation.

AI applications bring benefits including enhanced analysis of safeguards-relevant information and more efficient review of surveillance footage and video data, resulting in more effective use of scarce inspector resources (IAEA, 2023). Machine learning algorithms can process satellite imagery to detect changes at nuclear facilities, analyse trade data for indicators of undeclared nuclear activities, and screen scientific publications for proliferation-relevant research. The IAEA has been integrating these capabilities into its routine operations, with AI tools assisting analysts in processing the growing volume of open-source information that must be evaluated as part of the safeguards process.

The IAEA held its third Emerging Technologies Workshop in January 2025, bringing together experts from across the nuclear verification community to examine advancements in AI and explore their potential for enhancing nuclear safeguards. These workshops reflect a recognition within the IAEA that technological innovation is essential to maintaining the effectiveness of the safeguards system in an era of increasing demands and constrained resources.

Satellite imagery analysis powered by deep learning has shown particular promise for non-proliferation and counter-proliferation efforts. Studies have examined the applicability of deep-learning techniques for extracting artificial structures from high-resolution satellite imagery to support verification processes, enabling analysts to detect facility modifications, new construction, or other indicators of proliferation activity that might otherwise go unnoticed (Han et al., 2025). Deep learning models can be trained to recognise specific types of facilities, equipment, or operational patterns associated with nuclear fuel cycle activities, providing continuous monitoring capability that supplements traditional on-site inspection activities.

A 2024 technical report from the Oak Ridge National Laboratory provides a comprehensive assessment of how the IAEA is already utilising AI and machine learning to analyse open-source information for nuclear safeguards verification. The report documents that AI and machine learning tools already assist the IAEA with automating processes and analysing large datasets, including satellite imagery and unstructured data, thereby improving the

efficiency and effectiveness of safeguards implementation (Parada Iturria et al., 2024). This represents one of the most concrete examples of AI being deployed in an operational non-proliferation context by an international organisation.

AI excels at analysing dynamic and complex datasets, making it particularly effective for validating or physical claims about infrastructure changes at nuclear facilities. This capability is especially valuable given the volume of safeguards-relevant data that the IAEA and national authorities must process. However, the ORNL report also highlights important challenges, including the need for large, unbiased training datasets, the risk of AI-generated false information, and the potential for deliberate manipulation of satellite imagery to deceive AI systems. These challenges underscore the importance of maintaining human oversight and developing robust validation mechanisms for AI-assisted verification.

2.2 Chemical Verification

The Organisation for the Prohibition of Chemical Weapons (OPCW), the implementing body for the Chemical Weapons Convention (CWC), has launched multiple initiatives to explore the potential of AI for chemical non-proliferation and verification. The CWC, which entered into force in 1997, is the most successful disarmament treaty in history, having overseen the verified destruction of all declared chemical weapons stockpiles—totalling more than 72,000 metric tonnes of chemical agents. Maintaining this achievement and preventing the re-emergence of chemical weapons requires continuous innovation in verification and monitoring.

The OPCW has established a Scientific Advisory Board (SAB) Temporary Working Group on AI, which began its mandate on 1 January 2025, to provide expert advice on the implications of AI for the implementation of the CWC (OPCW, 2025a). This group is tasked with examining both the opportunities that AI presents for strengthening the Convention and the risks that AI-enabled technologies may pose to its object and purpose.

In June 2025, the OPCW held the Shanghai Workshop on Artificial Intelligence and Chemical Safety and Security Management—its first AI-focused capacity-building programme. The workshop brought together experts from across regions to discuss how AI can enhance chemical safety and security. The workshop explored topics including AI-powered monitoring of chemical production facilities, machine learning for detecting anomalies in chemical trade data, and the use of AI to strengthen national implementation of the CWC (OPCW, 2025b).

Perhaps the most significant OPCW initiative in this area is the AI Research Challenge, launched in July 2024 and funded by the European Union and the United Kingdom. In April 2025, the OPCW announced the selection of four research institutions to participate in the challenge: the University of Alberta (Canada) is developing AI-powered chemical language models to predict novel toxic compounds, creating a reference library that will improve the OPCW's ability to identify and monitor both known and unknown chemical warfare agents; additional projects address AI applications for chemical forensics, detection, and verification (OPCW, 2025c). These projects represent the first systematic effort by an international disarmament organisation to harness AI for operational verification purposes.

Civil society organisations and academic researchers have also called for AI-driven technologies to be incorporated into OPCW strategies to enhance detection, anomaly identification, and real-time monitoring of chemical activities (RAND Corporation, 2025). AI systems can analyse chemical production data, monitor precursor chemical supply chains, and flag suspicious procurement patterns that might indicate diversion from legitimate civilian use to prohibited purposes. Given the vast scale of global chemical production and trade, manual analysis of all potentially relevant data is simply not feasible—AI offers the only realistic path to comprehensive monitoring.

2.3 Biological Monitoring

In the biological domain, AI is reshaping biodefence and biosurveillance by improving detection capabilities and supporting compliance monitoring. The Biological Weapons Convention (BWC), unlike the NPT and CWC, lacks a formal verification mechanism. This gap has long been recognised as a weakness in the international non-proliferation architecture, and AI may offer new approaches to addressing it.

Modern biodefence requires enhanced biosurveillance, precision bioinformatics, and AI-driven threat detection to shift from reactive response to proactive prevention (Smith et al., 2026). The volume of biological data now being generated—from genome sequencing, proteomics, and other high-throughput techniques—far exceeds the capacity of human analysts to process. AI systems trained on biological datasets can identify patterns indicative of deliberate pathogen development, monitor the publication and patent literature for dual-use research, and support efforts to verify compliance with the BWC.

Researchers at the National Defense University have examined how AI could support enforcement of the Biological Weapons Convention in the absence of a formal verification mechanism. Their analysis assesses the capabilities and limitations of current AI systems for monitoring dual-use research, procurement networks, satellite imagery, and open-source data related to biological weapons programmes. They identify six core competencies necessary for AI-enabled verification: pattern recognition, anomaly detection, predictive analysis, natural language processing, image analysis, and data integration. At the same time, they emphasise that effective treaty enforcement requires purpose-built systems, transparency, auditability, and sustained human oversight—conditions that current AI systems do not yet fully satisfy (Annett et al., 2026).

The National Academies of Sciences, Engineering, and Medicine (2025) report on AI in the life sciences provides the most authoritative assessment to date of the biosecurity implications of advanced AI. The report acknowledges that AI tools, particularly large language models and AI-enabled biological design tools, have the potential to lower barriers to the misuse of biology, including by non-State actors. However, it also emphasises that the same tools can enhance biosecurity monitoring, improve screening of DNA synthesis orders, and strengthen the overall biosecurity ecosystem. The report recommends the development of governance frameworks that enable the beneficial applications of AI in the life sciences while mitigating risks of misuse—a balancing act that directly parallels the challenge of implementing UNSCR 1540.

AI-enabled biomedical data platforms are being developed as key assets for identifying and accelerating protective and therapeutic responses to emerging biological threats. These platforms can analyse genetic sequence data, monitor pathogen evolution, and detect early signs of deliberate biological agent development. The International Gene Synthesis Consortium (IGSC), an industry-led group of DNA synthesis companies, has established screening protocols that could be enhanced by AI-based sequence analysis tools, potentially enabling more accurate identification of DNA sequences of concern while reducing false positives that can impede legitimate research.

3. Improving Prediction, Risk Assessment, and Early Warning

Beyond verification of past or present activities, AI opens new possibilities for predictive capabilities in WMD threat assessment. The ability to anticipate proliferation risks before they materialise is a critical complement to traditional monitoring approaches. While verification tells us what has happened or is happening, prediction offers the possibility of intervening before a threat fully emerges.

AI can enhance preparation and prediction by modelling complex crises, rapidly iterating potential scenarios, and identifying non-traditional signatures and indicators that human analysts might overlook (U.S. Department of State, 2023). Machine learning systems trained on diverse datasets—including trade data, scientific publications, shipping records, and satellite imagery—can detect patterns consistent with proliferation activity at a scale and speed that would be impossible through manual analysis alone. These systems can flag emerging risks for human investigation, enabling analysts to focus their attention on the most promising leads.

One of the most promising applications of AI for early warning is in the analysis of procurement patterns. Proliferation networks must acquire specialised materials, equipment, and technology to develop WMD capabilities, and these procurement activities often leave traces in trade data, shipping records, and financial transactions. AI systems trained on known proliferation procurement patterns can identify suspicious transactions that may indicate a nascent WMD programme. This capability is particularly relevant to UNSCR 1540, which requires States to establish and maintain effective export and border controls—not least because such procurement networks may be operated by, or on behalf of, non-State actors seeking to acquire weapons of mass destruction. The same pattern-learning logic can therefore be directed at detecting the financing, brokering, and logistics of non-State proliferation networks as much as at State-level programmes.

The Defense Advanced Research Projects Agency's (DARPA) SIGMA programme has developed a persistent, real-time detection and early warning system for chemical and biological threats. The programme integrates networked sensors with AI-driven data analysis to detect trace chemical signatures across large geographic areas, providing continuous monitoring capability that strengthens overall situational awareness (DARPA, 2024). The follow-on SIGMA+ programme has expanded this approach to include biological threat detection, creating an integrated capability for chemical and biological early warning.

AI can also help reduce the risk of inadvertent escalation that can arise from misperception or miscommunication. By improving the accuracy and timeliness of threat assessment, AI systems can contribute to strategic stability—particularly when applied to monitoring and early warning rather than to automated decision-making in combat contexts. This principle has direct relevance to non-proliferation: accurate threat assessment reduces the likelihood of overreaction to ambiguous indicators, while also reducing the risk that genuine proliferation activities will go undetected.

Predictive models for proliferation risk are still at a relatively early stage of development compared to applications in verification and monitoring. The complexity of proliferation pathways, the diversity of actors involved—ranging from States to non-State procurement networks and individual technical specialists, and the intentional efforts of proliferators to conceal their activities make prediction inherently challenging. However, as AI systems become more sophisticated and training datasets become more comprehensive, the potential for meaningful prediction is likely to increase.

4. Strengthening Export Controls and Supply Chain Security

Export controls are a central pillar of UNSCR 1540 implementation. Operative paragraph 3 of the resolution requires States to establish and enforce controls on the export, transit, trans-shipment, and re-export of nuclear, chemical, and biological weapons-related materials, equipment, and technology. These controls are essential to preventing non-State actors from acquiring the means to develop WMD capabilities. Yet the scale and complexity of global supply chains make effective export control implementation extraordinarily challenging.

AI offers powerful tools for strengthening these controls in an era of increasingly complex global supply chains. The global trading system encompasses billions of individual transactions each year, and national export control authorities must screen a significant fraction of these transactions for potential proliferation relevance. Manual review of all potentially relevant transactions is simply not feasible; AI-driven screening tools offer the only realistic approach to comprehensive export control enforcement.

AI is itself a dual-use technology—driving commercial innovation while also enabling capabilities relevant to WMD development. The Bureau of Industry and Security (BIS) of the U.S. Department of Commerce issued a policy statement in May 2025 clarifying that certain activities involving advanced computing integrated circuits and related commodities used to train AI models may require export authorisation under the Export Administration Regulations. The BIS determined that access to advanced computing integrated circuits for training AI models has the potential to enable military-intelligence and WMD end uses, explicitly recognising the dual-use character of AI technology and the need to control its proliferation (Bureau of Industry and Security, 2025).

This same duality means that AI can be deployed to monitor supply chains against diversion to prohibited purposes. AI-driven analytics of publicly available procurement data can identify risks within chemical, biological, radiological, and nuclear material supply chains, pinpointing vulnerabilities such as obsolete security measures or unusual transaction patterns (Defense Innovation Unit, 2025). Machine learning models can be trained to recognise the procurement signatures of proliferation networks—patterns of purchasing, shipping, and financing that differ from legitimate commercial activity—and flag these patterns for human investigation.

AI algorithms can play a significant role in verifying the prohibition on development, production, acquisition, transfer, stockpiling, and use of biological weapons. They can help determine whether particular indicators emerging from available datasets warrant further investigation—a crucial capability given the enormous volume of data that export control authorities must process (Annett et al., 2026). Similar approaches have been applied to achieve unprecedented supply chain transparency for chemical, biological, radiological, and nuclear defence agencies.

Natural language processing and machine learning tools can screen scientific publications, patent filings, and equipment catalogues for dual-use items that may be subject to control. These tools can help national authorities keep pace with the rapid evolution of technology and ensure that control lists remain relevant and effective. As new technologies emerge and existing technologies find new applications, control lists must be continually updated—a task that AI tools can support by identifying relevant developments in the global scientific and technical literature. As UNICRI has emphasised, the wide reach of AI poses significant practical challenges for national authorities responsible for implementing UNSCR 1540, particularly in relation to strategic trade controls, where automated screening must be complemented by human assessment of dual-use technologies (Kim, 2024).

The challenge of export controls is particularly acute for AI-related technologies themselves. The same AI models and computing hardware that power beneficial applications in healthcare, energy, and communications can also be applied to WMD-relevant research. Export control authorities must therefore navigate a complex terrain, ensuring that AI technologies are available for peaceful purposes while preventing their diversion to prohibited WMD programmes. This challenge directly engages the principle articulated in UNSCR 1540 operative paragraph 9, which affirms that the resolution does not conflict with the rights to peaceful use of technology under the NPT, CWC, and BWC.

The Australia Group, an informal arrangement of States that coordinates export controls on chemical and biological weapons-related materials, has begun to consider the implications of AI for its control lists and guidelines. Similarly, the Nuclear Suppliers Group and the Missile Technology Control Regime are engaging with the challenges posed by AI-enabled technologies. These multilateral export control arrangements, which complement the binding obligations of UNSCR 1540, are essential forums for developing coordinated approaches to AI-related export controls.

5. Advancing International Cooperation and Capacity Building

The effectiveness of UNSCR 1540 depends not only on national implementation but also on international cooperation and capacity-building, particularly for States with limited resources. The 1540 Committee has consistently emphasised the importance of assistance and cooperation in its work, maintaining a matchmaking mechanism that connects States in need of assistance with States and international organisations that can provide it. AI can support these efforts in several important ways. The role of regional arrangements is equally significant: the ASEAN Regional Forum, for instance, has long championed the implementation of UNSCR 1540 alongside broader non-proliferation and export-control cooperation across the Asia-Pacific (ASEAN Regional Forum, 2004). Drawing on the 1540 Committee's guidance and its Matrix, which systematise assistance and cooperation, helps ensure that AI-enabled capacity-building is directed toward the priorities identified by States themselves (1540 Committee, 2017).

The UNSCR 1540 TechHorizon initiative, launched in partnership with the UN Office for Disarmament Affairs, brings together AI researchers, defence educators, export control experts, and disarmament professionals to explore how academic and technical communities can contribute to the implementation of the resolution. This initiative reflects a growing recognition of the need to bridge the gap between technology developers and policy practitioners in the non-proliferation domain (1540 Compass, 2025). By creating a platform for dialogue and collaboration across disciplinary boundaries, TechHorizon aims to ensure that AI applications for non-proliferation are developed with a clear understanding of the operational requirements of the 1540 regime.

The fourth issue of the 1540 Compass e-journal, published in May 2025, is dedicated entirely to the theme of emerging technologies, including AI. It explores how innovation can serve as both a challenge and a catalyst for stronger implementation of UNSCR 1540. Contributors analyse both the risks and the promise that AI and related technologies bring to non-proliferation, addressing topics such as the role of AI and non-State actors in WMD proliferation, how technologies can enable risk mitigation and strengthen compliance, and the importance of partnerships between academia and private industry in reinforcing implementation efforts (1540 Compass, 2025). The journal serves as a platform for knowledge-sharing and awareness-building, contributing to the broader goal of strengthening the global non-proliferation culture.

Under Canada's 2025 G7 presidency, the Global Partnership Against the Spread of Weapons and Materials of Mass Destruction has been developing a strategy to address WMD-relevant technologies. The Global Partnership, launched at the 2002 G8 Kananaskis Summit, is a multi-donor initiative that funds non-proliferation projects in States with limited resources. Its growing engagement with AI reflects a recognition that emerging technologies must be addressed as part of comprehensive non-proliferation capacity-building. The Global Partnership recognises that AI is already being employed for predictive maintenance and safety improvements across the nuclear material production lifecycle, and that these applications can be extended to non-proliferation purposes.

International efforts to govern AI in security contexts are also advancing through multilateral channels. The UN General Assembly has demonstrated broad international support for maintaining human control over security policy decisions involving AI, with significant majorities voting in favour of resolutions addressing AI risks in nuclear weapons contexts. This reflects an international consensus that while AI can augment human decision-making, it should not replace human judgment in matters affecting international peace and security. This principle has direct relevance to UNSCR 1540 implementation: AI tools can support verification, monitoring, and export control decisions, but the ultimate responsibility for compliance determinations must remain with human authorities. Comparative analyses of national and institutional AI principles, such as those by Anand and Deng (2023), further illustrate how States are working to reconcile the beneficial use of AI with responsibility and oversight, offering useful reference points for the non-proliferation community (Anand & Deng, 2023).

Comparative case studies of existing dual-use control regimes similarly provide lessons for the design of international AI governance institutions, underscoring the value of robust verification, transparent data-sharing, and adaptive mechanisms that can keep pace with rapid technological change (Wasil et al., 2024).

Multilateral approaches to AI governance, with the United Nations at their centre, are attracting growing international consensus as a means of strengthening risk prevention and control in AI. These developments reflect a broader acknowledgement that emerging technologies must be governed cooperatively if their security implications are to be managed effectively.

Capacity-building efforts are a particularly important area for AI applications. Many States, particularly developing countries, lack the technical infrastructure and expertise needed to fully implement their UNSCR 1540 obligations. AI-powered tools can help close these capability gaps: machine translation systems can make technical guidance available in multiple languages; AI-enhanced training platforms can provide tailored capacity-building programmes; and automated screening tools can assist national authorities with limited resources in implementing effective export controls. The OPCW's Shanghai Workshop and AI Research Challenge demonstrate how international organisations can support capacity-building in AI applications for non-proliferation, ensuring that the benefits of these technologies are accessible to all States Parties.

Beyond national implementation, a growing ecosystem of regional organisations, academic networks and professional bodies is helping to connect technical expertise with the non-proliferation community in ways that support awareness-raising, education and implementation-support. Regional security forums, knowledge-exchange platforms and research consortia each contribute to this ecosystem in complementary ways, providing neutral spaces in which researchers, practitioners and educators can compare experiences, share good practices and identify possible areas for knowledge exchange without prescribing particular courses of action to any State. The value of such networks lies partly in their diversity: by bringing together institutions and experts from different regions and disciplines, they help ensure that discussions of emerging technologies and dual-use governance reflect a broad range of perspectives rather than any single national or institutional vantage point.

Universities and research institutions have a distinctive contribution to make within this ecosystem. Beyond generating the underlying science, they serve as centres for the education of the next generation of scientists, engineers and security professionals, and as hubs for interdisciplinary dialogue between technical and policy communities. Research-security offices, institutional biosafety and dual-use review processes, and academic codes of responsible conduct all represent practical mechanisms through which universities can reinforce the objectives of UNSCR 1540 at the institutional level. These are awareness-building and good-practice measures rather than regulatory impositions; they operate through education, voluntary guidance and internal governance, and they can be adapted to local circumstances in fields as varied as artificial intelligence, the life sciences, materials science and chemical research. In this way, academic institutions complement the work of international organisations and national authorities by embedding an understanding of dual-use risk into the day-to-day practice of research and education.

Cooperation aimed specifically at the non-State-actor dimension of UNSCR 1540 deserves particular emphasis, because it is that dimension that most directly distinguishes the resolution from other elements of the non-proliferation regime. Assistance and knowledge-exchange efforts directed at procurement networks, brokering, financial transactions and the technical expertise that non-State actors might draw upon can help States strengthen their screening, monitoring and reporting capacities in ways that are consistent with existing obligations. AI-enabled tools for anomaly detection, transaction screening and open-source analysis are relevant here, and international training programmes, expert exchanges and twinning arrangements between institutions can help translate these capabilities into practice in resource-constrained settings. The aim is not to direct States towards particular policies, but to expand the range of options, good practices and collaborative arrangements available to them as they pursue the common objective of preventing the acquisition of weapons of mass destruction by non-State actors.

Taken together, these institutions and networks illustrate how research, education and cooperation can reinforce the objectives of UNSCR 1540 through voluntary, awareness-based and implementation-support mechanisms. They embody the future collaboration pathways and good practices that the international community can continue to develop as emerging technologies evolve, providing a durable foundation on which to build sustained engagement between academic, technical, industry and policy communities. As the relationship between artificial intelligence and non-proliferation becomes better understood, such networks are well placed to capture emerging lessons, update educational materials and disseminate knowledge more widely, helping to ensure that awareness keeps pace with technological change.

6. Detecting Chemical and Biological Threats

The detection of chemical and biological threats presents unique challenges that AI is particularly well-suited to address. Chemical and biological agents are inherently difficult to monitor—they can be produced in small quantities using dual-use equipment, and their legitimate applications in medicine, agriculture, and industry create significant ambiguity about intent. Distinguishing between legitimate civilian activities and prohibited weapons development requires sophisticated analytical capabilities that AI can help provide.

In the chemical domain, AI systems can analyse spectral data from environmental sensors to detect trace quantities of chemical warfare agents or their precursors. Machine learning models trained on chemical signatures can distinguish between legitimate industrial emissions and signs of prohibited activities, reducing false alarms while maintaining high detection sensitivity (RAND Corporation, 2025). This capability is particularly valuable for wide-area monitoring: networked sensors coupled with AI-driven data analysis can provide continuous surveillance coverage over large geographic areas, detecting chemical threats that might otherwise go unnoticed until they have already caused harm.

The integration of AI with networked sensor systems creates the potential for wide-area monitoring that was previously impractical. DARPA's SIGMA programme has demonstrated this approach at scale, deploying networked chemical sensors across urban areas and analysing the resulting data stream with AI algorithms to detect trace chemical signatures. The programme has shown that AI can enable persistent, real-time monitoring for chemical threats in complex environments, providing early warning capabilities that were previously unavailable.

In the biological domain, AI-driven bioinformatics tools can screen synthetic DNA sequences for similarities to known pathogens or toxin genes, flagging potentially dangerous orders before they are fulfilled. This capability is becoming increasingly important as DNA synthesis technology becomes more accessible and the barrier to ordering custom genetic material continues to fall. Concerns about the potential for AI-enabled biological design tools to lower barriers to the creation of novel pathogens have been a focus of policy attention. The National Academies (2025) report notes that while current AI tools do not yet enable the design of novel pathogens from scratch, the rapid pace of advancement in this area warrants vigilance and the development of appropriate governance mechanisms.

Several commercial DNA synthesis companies have adopted AI-based screening protocols as part of their biosecurity measures. The International Gene Synthesis Consortium (IGSC), whose members account for the majority of global DNA synthesis capacity, has established harmonised screening standards that include sequence screening against databases of pathogens and toxins. AI-enhanced screening tools could improve the accuracy and efficiency of these protocols, enabling more precise identification of sequences of concern while reducing the burden on legitimate customers.

AI systems can also support epidemiological surveillance by analysing global health data, scientific literature, and other information sources for early indicators of unusual disease outbreaks—whether natural, accidental, or deliberate. Machine learning models trained on historical outbreak data can identify anomalies in disease patterns

that may warrant investigation. The Federation of American Scientists (2026) has examined the potential of AI for biological threat detection, highlighting both the promise of these approaches and the challenges of calibration and validation.

The dual-use nature of AI in the chemical and biological domains requires careful consideration. The same AI tools that can detect threats can also, in the wrong hands, accelerate the design of novel chemical agents or pathogens. The RAND Corporation (2025) report on mitigating risks at the intersection of AI and chemical and biological weapons provides a comprehensive assessment of this duality. It identifies several categories of risk, including the potential for AI to enable the discovery of novel chemical warfare agents, to optimise biological weapon development pathways, and to help malicious actors bypass existing detection and control measures.

This risk underscores the importance of governance frameworks that promote responsible use while preventing misuse. The OPCW's AI Research Challenge offers a model for how international organisations can harness AI for verification purposes while managing dual-use risks: by directing AI development toward specific, agreed-upon verification applications, and by ensuring that the resulting tools are deployed under institutional oversight rather than in uncontrolled environments. Similar approaches could be adopted in the biological domain, where AI tools developed for biosecurity and biosurveillance purposes could be governed through mechanisms that ensure their use remains within agreed parameters.

7. Academic Analysis and Relevance to UNSCR 1540

UNSCR 1540 establishes obligations directly relevant to the development and deployment of AI applications for non-proliferation. The resolution's operative paragraphs create a framework that both enables and constrains the use of AI for these purposes. Understanding this framework is essential for researchers, practitioners, and policymakers seeking to harness AI for the objectives of the resolution.

Table 1 provides a structured mapping of selected UNSCR 1540 operative paragraphs to specific AI capabilities and real-world deployment examples. References are drawn from authoritative government reports, international organisation documents, and peer-reviewed research.

Taken together, the mapping in Table 1 demonstrates that AI applications are relevant to nearly every aspect of UNSCR 1540 implementation — from the core obligation to prevent support to non-State actors (OP 1), through the detailed requirements for domestic controls (OP 2 and 3), to the mechanisms for international cooperation and capacity-building (OP 6, 7, and 8).

However, the table also reveals important gaps and challenges. Several cells represent potential rather than current reality: AI applications for financial screening related to proliferation financing, for example, are less developed than applications in verification and monitoring. The governance mechanisms called for in OP 9 are still at an early stage of development for AI technologies, and significant work remains to be done to ensure that the benefits of AI are accessible to all States while the risks are effectively managed.

The dual-use character of AI creates a particular challenge for the implementation of UNSCR 1540. The same AI capabilities that can strengthen non-proliferation can also, if misused, support proliferation activities. This duality mirrors the broader challenge of implementing UNSCR 1540 in an era of rapid technological change, and it underscores the importance of the governance frameworks that the resolution establishes.

8. Implications for Awareness, Education, Research Engagement, and Implementation-Support

The integration of AI into non-proliferation efforts carries significant implications for awareness-building, education, research engagement, and implementation-support activities. These implications are relevant to the full range of stakeholders involved in UNSCR 1540 implementation, including national governments, international organisations, academic institutions, civil society, and the private sector.

Awareness-building efforts must address both the opportunities and the risks that AI presents for non-proliferation. Policymakers, industry professionals, and the broader public need to understand how AI can support the objectives of UNSCR 1540 while also recognising the risks that AI itself poses as a dual-use technology. The 1540 Compass e-journal, the TechHorizon initiative, and similar platforms play an important role in this awareness-building, providing accessible analysis of complex technical topics for non-specialist audiences.

Educational institutions have an important role to play in developing curricula that bridge the gap between technical AI training and non-proliferation studies. University programmes in international security, disarmament, and non-proliferation have traditionally focused on legal, political, and diplomatic dimensions of the field. As AI becomes increasingly relevant to non-proliferation practice, these programmes must evolve to include technical content that enables students to understand both the capabilities and the limitations of AI systems. Conversely, computer science and AI programmes should incorporate content on dual-use risks and responsible innovation, ensuring that the next generation of AI developers is aware of the security implications of their work.

Research engagement is essential to ensuring that AI applications for non-proliferation are developed with the specific requirements of the domain in mind. Academic researchers in computer science, international security, and related fields can contribute to the development of AI tools that are transparent, auditable, and appropriate for use in security contexts. The 1540 TechHorizon initiative and the OPCW AI Research Challenge offer models for how researchers and practitioners can collaborate effectively, directing AI development toward agreed non-proliferation objectives rather than leaving it to market forces alone.

The OPCW's AI Research Challenge is particularly instructive as a model for research engagement. By identifying specific verification and detection challenges that AI can address, and by funding targeted research projects to develop solutions, the OPCW has created a framework that harnesses academic and technical expertise for concrete non-proliferation purposes. Similar challenge-based approaches could be adopted in other domains, including nuclear safeguards verification and biological threat detection.

Implementation-support activities—including technical assistance, training, and capacity-building for States—can benefit significantly from AI-enabled tools. AI-powered translation and document analysis systems can help States prepare their national implementation reports for the 1540 Committee, reducing the burden of compliance reporting and improving the quality and consistency of submissions. Machine learning tools can assist national authorities in screening export licence applications and identifying suspicious transaction patterns, providing capabilities that would otherwise require substantial investments in human expertise. Online training platforms enhanced by AI can provide tailored capacity-building programmes for officials from States with limited resources, adapting content to the specific needs and circumstances of each user.

The dual-use risks of AI also have implications for awareness and education. The same AI capabilities that can be applied for non-proliferation purposes can also be misused for proliferation, and awareness of this dual-use character is essential for responsible development and deployment. Educational programmes that address the dual-use implications of AI are an important component of the broader effort to build a culture of responsibility in the AI research and development community.

Consistent with the publication's guiding principles, contributions are intended to frame analysis as academic work contributing to awareness, education, research engagement, or implementation-support, rather than as policy recommendations. This framing is appropriate for the present article: the goal is not to prescribe what governments should do, but to explore how AI can be understood and potentially applied within existing non-proliferation frameworks. The discussion of AI applications presented here is intended to inform academic debate, support educational activities, and contribute to a broader understanding of the relationship between emerging technologies and the objectives of UNSCR 1540.

8.1 Practical Examples and Implementation-Support in Practice

The preceding sections surveyed how AI can support civilian-led, non-State-actor-centred implementation of UNSCR 1540. This sub-section offers concrete, applied examples drawn from existing programmes to illustrate how awareness-building, education, research engagement, and implementation-support can be operationalised in practice. The examples are deliberately presented at a general, capability level rather than as detailed technical specifications, reflecting the broad categories of machine learning, deep learning, natural language processing, and computer vision that underpin this field.

In the nuclear domain, the IAEA already applies machine learning and computer vision to the analysis of safeguards data, including satellite imagery and unstructured open-source information, to help verify the completeness of State declarations and to flag anomalies that warrant further human review (Parada Iturria et al., 2024). Such tools do not replace inspector judgment; they triage and prioritise information so that scarce expert capacity is directed where it matters most. For States and the Secretariat alike, this offers a practical, low-cost path to stronger verification without new treaty obligations.

In the chemical domain, the OPCW's AI Research Challenge demonstrates how a challenge-based model can convert general technical interest into targeted, verifiable capability improvements. The four projects selected in the first cycle of the Challenge are designed to strengthen chemical identification and forensic analysis, and the effort is accompanied by an annual capacity-building programme focused on AI for chemical safety and security management (Organisation for the Prohibition of Chemical Weapons, 2025a). This combination of funded research and sustained training provides a replicable template for other technical bodies seeking to industrialise AI development for non-proliferation.

In the biological domain, the absence of a formal verification mechanism under the Biological Weapons Convention places particular weight on monitoring and information-sharing arrangements such as the International Gene Synthesis Consortium (IGSC). Sequence-screening tools that flag synthetic DNA requests matching dangerous pathogens are a concrete, immediately deployable application of pattern-matching techniques that States and private providers can integrate into existing supply-chain controls (Annett et al., 2026). Analysed alongside studies of laboratory oversight and reporting (Smith, Brockmann, & Bromley, 2026), these measures show how AI can reinforce governance functions that rely on voluntary cooperation rather than inspection.

For supply-chain and export-control authorities, natural language processing tools can assist national authorities in screening licence applications, patent filings, and equipment catalogues for dual-use indicators, reducing the burden on human analysts and improving consistency across submissions (Kim, 2024). On the physical-security side, the Defense Innovation Unit's partnership with commercial providers to apply AI-driven analytics to chemical, biological, radiological, and nuclear material supply chains illustrates how procurement transparency can be enhanced through automated early-warning of obsolescence, natural-disaster, or hostile-activity vulnerabilities (Defense Innovation Unit, 2025).

Finally, for States with limited technical infrastructure, AI-enabled machine translation and document-analysis tools can lower the cost of preparing national implementation reports for the 1540 Committee, while adaptive online training platforms can tailor capacity-building content to the specific needs of individual officials. These examples share a common practical logic: they use AI to amplify, rather than replace, human expertise, and they direct scarce resources toward the highest-impact verification, monitoring, and assistance functions that underpin the non-State-actor objectives of UNSCR 1540.

9. Conclusion

Artificial intelligence is not a panacea for the challenges of WMD non-proliferation. As numerous commentators have observed, AI is a double-edged sword: it can potentially camouflage proliferation activities and lower barriers to entry for malicious actors. The same machine learning tools that enable early detection of biological threats could, in the wrong hands, accelerate the design of novel pathogens. The same AI-driven analytics that strengthen export controls could be used to identify and exploit gaps in national implementation measures. The National Academies (2025) have warned that the rapid pace of AI advancement in the life sciences may outstrip the development of governance mechanisms, creating a window of vulnerability that could be exploited by malicious actors.

However, these risks do not diminish the potential of AI to make a meaningful contribution to the non-proliferation objectives of UNSCR 1540. With appropriate governance frameworks, sustained international cooperation, and a commitment to civilian applications, AI can take major risks off the table. AI systems can process data at scales that human analysts cannot match, detect patterns that would otherwise remain invisible, and automate routine monitoring tasks to free scarce human expertise for higher-level analysis. The IAEA is already using AI to enhance its safeguards activities, the OPCW has launched an ambitious AI Research Challenge, and the 1540 community is actively exploring the implications of emerging technologies through initiatives such as TechHorizon and the 1540 Compass.

The path forward requires a balanced approach—one that harnesses AI's capabilities for verification, monitoring, prediction, and supply chain security while maintaining meaningful human control over decisions that affect international peace and security. The international community, working through the UN, IAEA, OPCW, the 1540 Committee, and other multilateral institutions, must continue to develop the norms, capacities, and partnerships needed to ensure that AI serves as a shield against proliferation rather than a tool for it.

Several principles should guide this effort. First, human oversight must remain central to all AI applications in non-proliferation. AI tools can inform and support human decision-making, but the ultimate responsibility for compliance determinations must rest with human authorities. Second, transparency and auditability are essential for building trust in AI-assisted verification and monitoring. States must be able to understand how AI systems reach their conclusions and must have confidence that these systems are not subject to manipulation or bias. Third, international cooperation and capacity-building are essential to ensuring that all States, regardless of their resources, can benefit from AI applications for non-proliferation. The digital divide must not become a non-proliferation divide. Fourth, governance frameworks must evolve in parallel with the technology itself, anticipating risks before they materialise rather than responding after the fact.

Ultimately, AI is what we make of it. With deliberate effort, international cooperation, and a clear-eyed understanding of both its potential and its limitations, AI can become a valuable component of the global effort to prevent the proliferation of nuclear, chemical, and biological weapons to and by non-State actors—and a worthy complement to the vision of international security that UNSCR 1540 represents.

References

1540 Committee. (2017). 1540 Matrix template. United Nations Security Council. https://www.un.org/en/sc/1540/national-implementation/1540-matrices.shtml

1540 Compass. (2025). *Fourth issue of 1540 Compass: Emerging Technologies: Challenges and Opportunities*. United Nations Interregional Crime and Justice Research Institute (UNICRI). https://unicri.org/index.php/Publications-Magazine-1540%20Compass-4

Anand, A., & Deng, H. (2023, February 13). *Towards responsible AI in defence: A mapping and comparative analysis of AI principles adopted by States*. UNIDIR, Geneva, Switzerland.

Annett, E., Giordano, J., & Melley, B. (2026, February 20). How AI can help enforce the Biological Weapons Convention. National Defense University, Institute for National Strategic Studies, Center for Disruptive Technology and Future Warfare. https://digitalcommons.ndu.edu/cdtfw-articles/20/

ASEAN Regional Forum. (2004, July 2). *ASEAN Regional Forum Statement on Non-Proliferation* [Statement]. The 11th Meeting of the ASEAN Regional Forum, Jakarta, Indonesia.

Bureau of Industry and Security (2025). *Policy statement on controls that may apply to advanced computing integrated circuits and other commodities used to train AI models* (13 May 2025). U.S. Department of Commerce. https://www.bis.gov/media/documents/ai-policy-statement-training-ai-models-may-13-2025

Defense Advanced Research Projects Agency (2024). *SIGMA+ program overview*. https://www.darpa.mil/program/sigma

Defense Innovation Unit. (2025). *Resilinc — Supply Chain Illumination* [Solution profile]. DIU Product Catalog. https://www.diu.mil/solutions/portfolio/catalog/a0Tt0000009WqvtEAC-a0hcr000005DDpOAAW

Federation of American Scientists (2026). *Artificial intelligence and biological risks: Current status and future risks*. FAS Global Risk Report.

Han, J.-J., Ha, G., Han, Y., Lee, C., Lee, H., & Song, A. (2025). Deep learning applications on satellite imagery datasets for nuclear nonproliferation and counter-proliferation. *Annals of Nuclear Energy*, 219, 111443. https://doi.org/10.1016/j.anucene.2025.111443

International Atomic Energy Agency (2023). The evolution of safeguards technology. *IAEA Bulletin*. https://www.iaea.org/bulletin/the-evolution-of-safeguards-technology

Kim, H. (2024). Artificial intelligence: Emerging challenges in implementing UNSCR 1540. *1540 Compass*, (2). United Nations Interregional Crime and Justice Research Institute (UNICRI).

National Academies of Sciences, Engineering, and Medicine (2025). *The age of AI in the life sciences: Benefits and biosecurity considerations*. Washington, DC: The National Academies Press. https://doi.org/10.17226/28868

Organisation for the Prohibition of Chemical Weapons (2025a). *Note by the Technical Secretariat: Workshop on Artificial Intelligence and Chemical Safety and Security Management: Outcomes and Key Findings* (S/2444/2025, 17 September 2025).

Organisation for the Prohibition of Chemical Weapons (2025b). *Report on the Shanghai Workshop on Artificial Intelligence and Chemical Safety and Security Management*. OPCW Technical Secretariat.

Organisation for the Prohibition of Chemical Weapons (2025c). *OPCW AI Research Challenge: Harnessing AI tools to enhance global chemical security* [Press release, 29 April 2025]. https://www.opcw.org/media-centre/news/2025/04/opcw-ai-research-challenge-harnessing-ai-tools-enhance-global-chemical-security

Parada Iturria, F., Martindale, N., Reasor, A., Stewart, S., & Ukishima, L. (2024). *AI for nuclear safeguards verification*. Oak Ridge National Laboratory. https://doi.org/10.2172/3002744

Puscas, I. (2024, November 6). *Large language models and international security: A primer*. UNIDIR, Geneva, Switzerland.

RAND Corporation (2025). *Mitigating risks at the intersection of artificial intelligence and chemical and biological weapons* (RR-A2990-1). Homeland Security Operational Analysis Center. https://www.rand.org/pubs/research_reports/RRA2990-1.html

Reinhold, T., Hoffberger-Pippan, E., Blanchard, A., Blum, M.-M., Lentzos, F., & Saltini, A. (2025). *Artificial intelligence, non-proliferation and disarmament: A compendium on the state of the art* (EU Non-Proliferation and Disarmament Papers, No. 92). Stockholm International Peace Research Institute (SIPRI).

Smith, M., Brockmann, K., & Bromley, M. (2026). *Preventing biological weapons proliferation: Operational applications of emerging technologies*. Stockholm International Peace Research Institute. https://doi.org/10.55163/HPTG4304

U.S. Department of State (2023). *Report of the International Security Advisory Board on artificial intelligence and associated technologies*.

United Nations Security Council (2004). *Resolution 1540 (2004)* (S/RES/1540(2004)).

Wasil, A. R., Barnett, P., Gerovitch, M., Hauksson, R., Reed, T., & Miller, J. W. (2024). Governing dual-use technologies: Case studies of international security agreements and lessons for AI governance. *arXiv preprint 2409.02779*. https://doi.org/10.48550/arXiv.2409.02779

ARTICLE 04

AI-Enabled Export and Border-Control Awareness under UNSCR 1540: A Narrative Overview through Goods, Data and Routes

Yu Guan
Associate Professor
Department of Computer Science, University of Warwick

Sai Gu
Deputy Pro-Vice-Chancellor (China) and Professor of Chemical and Process Engineering
University of Warwick



Abstract

The rapid development of emerging technologies is reshaping the practical environment in which United Nations Security Council Resolution 1540 (UNSCR 1540) is understood and implemented. UNSCR 1540 aims to prevent non-State actors from acquiring nuclear, chemical and biological weapons, related materials and means of delivery. In practice, the implementation of UNSCR 1540 depends partly on the interpretation of goods, data, documents, routes, actors and services across borders. Traditional control lists are essential, but modern trade, digital knowledge transfer, service-based production and AI-mediated information access complicate the relationship between listed items and real-world risk assessment. This narrative overview discusses how AI-enabled methods such as vision-language models (VLMs), multimodal large language models (MLLMs) and remote sensing may contribute to awareness and implementation-support discussions around export control and border control under UNSCR 1540. Specifically, we adopt a goods–data–routes framework to examine how AI-assisted interpretation may help human experts identify and assess ambiguity across physical items, associated information and patterns of movement. We also briefly discuss potential chemical, biological, radiological and nuclear (CBRN) information risks, including how AI may lower certain informational barriers while significant

physical, technical and operational constraints remain. The main contribution of this article is to synthesise and translate selected AI research into language that supports UNSCR 1540 awareness, education, research engagement and implementation-support discussions. This could provide a starting point for further interdisciplinary dialogue on the responsible use of AI in export- and border-control contexts.

Keywords: UNSCR 1540; export control; border control; AI; vision-language models; multimodal large language models; remote sensing

Introduction

UNSCR 1540 is commonly understood through its legal and institutional vocabulary: non-State actors, nuclear, chemical and biological weapons, weapons of mass destruction (WMD), related materials, means of delivery, domestic controls and national implementation. Adopted under Chapter VII of the UN Charter, the resolution requires all States to adopt and enforce appropriate effective measures to prevent non-State actors from manufacturing, acquiring, possessing, developing, transporting, transferring or using such weapons and their means of delivery (United Nations Security Council, 2004). Conceptually, UNSCR 1540 can be understood as a universal prevention framework within the wider non-proliferation architecture. Its practical significance lies not only in prohibiting WMD-related activities by non-State actors, but also in strengthening the control environments through which sensitive materials, technologies, services and forms of know-how are secured, monitored and regulated. This article applies this prevention logic to the practical implementation of export and border controls in an emerging technology landscape.

AI-enabled systems may help make complex information easier to interpret in contexts where officials, companies, researchers and practitioners must reason across incomplete descriptions, customs codes, images, sensor readings, routes, end users and digital records. In this sense, AI is treated as part of the changing environment of UNSCR 1540 implementation-support discussions, helping technical researchers, policy communities and practitioners understand how export- and border-control problems can be shaped by data, sensing and interpretation. Awareness in this context includes understanding what AI can do, what it cannot do, where it may introduce new risks, and how it may support human analysis of ambiguous information.

There are several forms of ambiguity within export and border control. For example, a product description may be incomplete or misleading. An HS code may be too broad to express export-control sensitivity. A product image or X-ray image may not match the declared description. A transaction may involve a legitimate commercial service, while the final end-use remains unclear. A design file may move electronically rather than physically. A route may involve trans-shipment, re-export or an intermediary with limited visible connection to the final user. These are not purely technical problems, but they are also not purely legal problems. They require interpretation across multiple types of evidence.

In this article, we adopt a goods–data–routes framework to explain how AI-assisted interpretation may make ambiguity more visible to human experts. Specifically, we treat export and border control as multimodal interpretation problems. 'Multimodal' simply means that relevant evidence may come in different forms: text, images, sensor readings, shipping data, technical specifications, satellite imagery, company information and trade routes. A customs officer, export-control analyst, compliance officer or researcher may need to connect these signals to form a careful view of risk. AI-enabled methods, including vision-language models (VLMs), multimodal large language models (MLLMs) and remote-sensing methods, may help organise and interpret these signals, as illustrated in Fig. 1.

This article provides an interdisciplinary narrative overview rather than a systematic review or an evaluation of operational AI systems. It draws together selected official, academic and technical literature to explain how AI-enabled methods may relate to export- and border-control awareness under UNSCR 1540. The goods–data–routes framework is used as an organising device for this discussion rather than as a replacement for established legal or operational classifications.

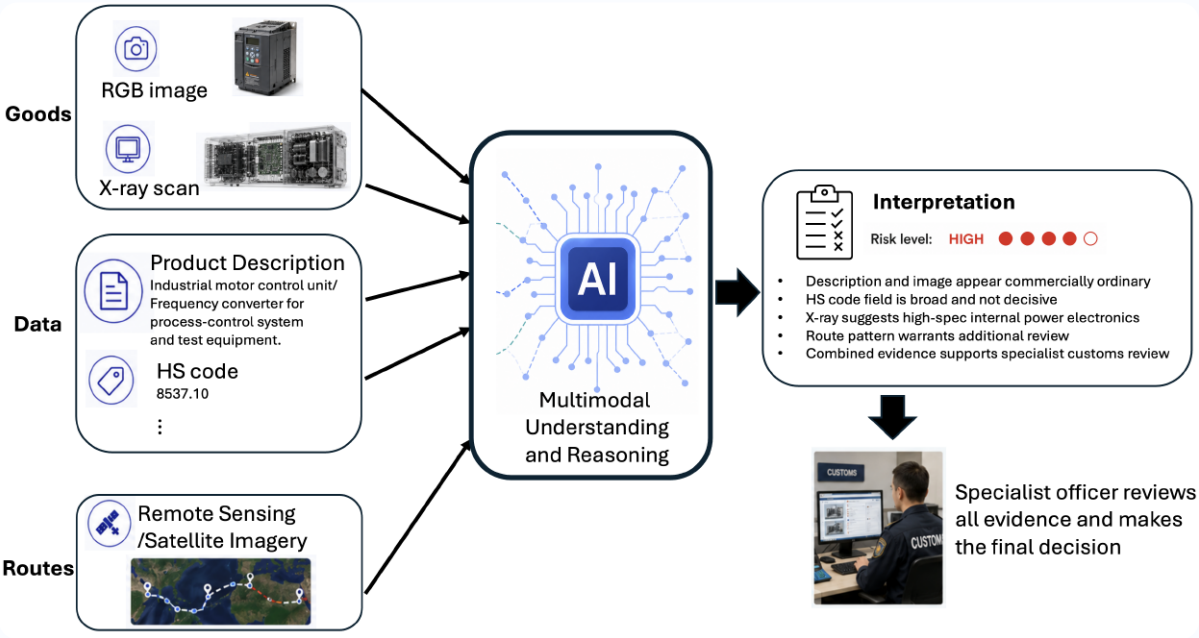

Figure 1. An illustrative AI-enabled workflow for interpreting ambiguity across goods, data and routes in export- and border-control contexts.

In this article, we also respond to a growing public discussion about AI and chemical, biological, radiological and nuclear (CBRN) risks. Claims that advanced AI models allow non-experts to develop CBRN weapons in a backyard are often too simplistic. Current evidence suggests that AI may lower certain informational and organisational barriers by supporting information retrieval, synthesis, translation, planning, troubleshooting and workflow integration, particularly in chemical and biological contexts. However, access to physical materials, specialised equipment, tacit expertise, suitable facilities, testing capabilities and supply chains remains subject to substantial practical and legal constraints (Mouton et al., 2023; OpenAI, 2024a). This distinction calls for proportionate caution and reinforces the continuing importance of appropriate and effective export and border controls over relevant materials, equipment, technologies and associated services in limiting the translation of AI-assisted information into practical CBRN capabilities.

Export and Border Controls under the UNSCR 1540 Framework

UNSCR 1540 was adopted in 2004 as part of the international response to the risk that non-State actors might acquire WMD, related materials or means of delivery. It establishes a broad obligation for all States to adopt and enforce effective measures against non-State actor proliferation involving nuclear, chemical and biological weapons and their means of delivery (United Nations Security Council, 2004). Specifically, operative paragraph 3 of UNSCR 1540 calls on States to develop and maintain appropriate effective measures for accounting for and securing relevant items, physical protection, border controls and law enforcement, and export, transit, trans-shipment and re-export controls, including controls on related funds and services (United Nations Interregional Crime and Justice Research Institute, 2024). These categories make clear that the resolution is not only about weapons as finished objects, but also about materials, technologies, services, movement, intermediaries and enabling networks.

This matters because the emerging technology landscape is not confined to laboratories or digital platforms. It is also embedded in supply chains. Additive manufacturing services, automated laboratories, cloud-based design tools, digital marketplaces, specialised logistics providers and remote analytical services can change how goods,

knowledge and capabilities move. UNSCR 1540 remains technology-neutral, but the practical environment in which export and border controls operate is increasingly reshaped by emerging technologies. The resolution's continued relevance depends partly on how its core categories are understood in this environment.

The term 'related materials' is important. It is sometimes used loosely as if it simply means dual-use goods, but the concept is broader. Related materials may include materials, equipment and technology covered by relevant treaties, multilateral arrangements or national control lists. Some items are dual-use in the familiar sense: they have ordinary civilian applications but may also contribute to WMD-related activities. Other items may be more specifically designed for sensitive purposes. For awareness and education, it is useful to preserve this broader understanding rather than reducing the discussion to dual-use technology alone.

Export control and border control are practical sites where 1540-related concerns become visible. Export control concerns the movement of goods, software, technology and services out of a jurisdiction or onward to another user. Border control concerns what crosses land, maritime and air borders, while export-control processes may also address certain digital transfers. Trans-shipment and re-export controls concern goods that pass through or move onward. These processes bring together legal classification, customs practice, technical expertise, documentation, intelligence, compliance and enforcement. Traditional export-control systems rely heavily on control lists. National lists may draw on those developed through multilateral export-control arrangements, including the Nuclear Suppliers Group, the Australia Group, the Missile Technology Control Regime, the Wassenaar Arrangement and the Zangger Committee (Stockholm International Peace Research Institute, 2002; United Nations Security Council 1540 Committee, n.d.). Such lists help identify known sensitive materials, equipment, software, technology and means of delivery. They provide legal clarity and a common vocabulary for licensing, compliance and enforcement.

At the same time, list-based systems face practical interpretation challenges. A list entry may require technical expertise to apply. For example, a commodity code may not map neatly onto a control-list category; a product description may omit critical performance characteristics; a shipment may be composed of ordinary components that only become concerning in combination; a digital file may carry design knowledge without a physical item crossing a border. These challenges are not defects in control lists as such; they are features of the complex environment in which lists must be interpreted.

These implementation challenges reinforce the importance of assistance, outreach, cooperation and capacity-building within the UNSCR 1540 framework. Emerging-technology research can contribute by clarifying what information is available, how evidence from different sources may be combined, where uncertainty arises and why human judgement remains essential. As export- and border-control capabilities vary across jurisdictions and institutional settings, this article presents AI-enabled systems as tools for awareness, interpretation, education and research engagement rather than as universal technical solutions.

From Control Lists to Ambiguity

For analytical purposes, this article organises export- and border-control information into three interconnected domains: goods, data and routes. Goods comprise physical materials, components and equipment, which may be represented in screening systems through product photographs, X-ray scans or other images. Data include declarations, technical specifications, HS codes, digital files and transaction records. Routes describe the geographical and organisational pathways through which goods and information move.

Export and border-control ambiguity often emerges because these three categories do not align neatly. A goods-only view can miss contextual risk. Many goods with potential WMD relevance also have legitimate civilian uses. For example, pumps, valves, machine tools, sensors, fermenters, drones, composite materials, laboratory equipment and electronic components can have legitimate commercial applications. Their relevance may depend on performance thresholds, material properties, precision, end-use, combination with other goods or the identity of the consignee. A data-only view can also be misleading because customs documents and HS codes serve broad

administrative and commercial purposes and do not correspond directly to export-control classifications. For example, a single HS category may include both sensitive and non-sensitive goods, creating ambiguity when assessing strategic relevance.

Digital transfers further complicate the relationship between goods and data. Design files, technical drawings, algorithms, manufacturing instructions and biological sequence information may be transmitted electronically without a corresponding physical shipment. Such information can be copied, modified and transferred rapidly, creating practical challenges for export-control and research-security processes. Service-based production further blurs this relationship. For example, users may submit digital designs to additive-manufacturing providers, outsource synthesis or testing to specialised laboratories, or request simulation and optimisation services. Although these activities commonly serve legitimate purposes, they can separate the requester, the technical information and the physical production capability, making end-use assessment and due diligence more difficult.

Routes introduce a separate source of contextual ambiguity. Legitimate trade routes can be complex for ordinary commercial reasons: cost, availability, warehousing, shipping schedules, free zones, regional distribution and re-export markets. Proliferation networks may exploit complexity by using front companies, intermediaries, trans-shipment hubs or fragmented procurement. Route information does not automatically imply risk, but unusual patterns can provide context when combined with goods and data.

Taken together, these examples do not suggest that ambiguity can be eliminated through item classification or automated analysis alone. Rather, they show why assessment may require contextual information about the goods, supporting data, service relationships, actors, intended end use and routes involved. The role considered for AI in this article is therefore supportive: AI-enabled systems may help organise these signals, identify inconsistencies and direct matters to expert review, while legal interpretation and operational decisions remain human responsibilities.

The concept of ambiguity is not a claim that control systems are failing. It reflects the complexity of assessing export- and border-control information from partial and sometimes inconsistent signals. Export and border-control actors often work with partial signals. They may have a declared description, an HS code, a product image, an X-ray scan, a shipping route, a consignee name and a limited time window. The question is how these signals can be interpreted together, and how uncertainty can be communicated responsibly.

AI-enabled systems become relevant because they are increasingly able to process multiple forms of data. A language model can summarise documents or compare product descriptions. A vision model can analyse images. A vision-language model (VLM) can connect images and text. A multimodal system can combine several evidence streams and a remote sensing model can provide information about routes, facilities, ports or wide-area movement patterns by analysing satellite or aerial images. Such AI-enabled systems may support the assessment of ambiguous cases by highlighting mismatches, organising evidence, prioritising records for review, identifying patterns or explaining why a case appears unusual.

AI Models and AI-Enabled Interpretation

This overview considers three broad categories of AI-enabled interpretation: vision-language models (VLMs), multimodal large language models (MLLMs) and remote-sensing methods.

VLMs are AI models that can look at images and text together, and they became widely used after research demonstrated that models could learn joint representations of images and text at large scale (Radford et al., 2021). In practical terms, such models can associate visual patterns with textual descriptions and can support tasks such as image captioning, visual question answering, retrieval and mismatch detection. More technical details about VLMs can be found in Appendix. In this context, a VLM may compare a product image with a declared description and assess whether the two appear consistent. That is, it does not need to know that a particular item is controlled in order to notice that a declared 'plastic household component' does not visually resemble the object in the image.

This image–text cross-checking capability may be useful in export- and border-control scenarios, especially when product descriptions are incomplete, generic or misleading. For instance, a product image may reveal features not mentioned in the description; an X-ray scan may show an internal structure that differs from the declared content. In such cases, VLMs may help flag inconsistencies for human review.

A multimodal large language model (MLLM) can process or organise more than one type of input. These inputs may include RGB images, X-ray scans, radiation readings, chemical sensor data, shipping records, documents and route information. Border-control environments are already multimodal in this sense. A container may be associated with a manifest, an X-ray scan, a seal number, a route, a consignee and a risk score. The technical question is how these evidence streams can be organised and interpreted together in ways that remain intelligible to human users. MLLMs commonly use modality-specific encoders to transform heterogeneous inputs into numerical representations, which are then mapped into a form that a language model can process for tasks such as question answering or summarisation (OpenAI, 2024b). More technical details about MLLMs can be found in Appendix. In this context, an MLLM may help compare declarations with images, align sensor readings with shipment records, associate route changes with documentation, or present a summary of why a case appears unusual, as shown in Fig. 1. MLLMs may also help non-specialists access technical information by translating complex data into structured prompts for expert review. This may make fragmented evidence easier to examine together.

Remote sensing adds a spatial and route-oriented dimension. Remote sensing generally means observing places or movement from a distance, commonly through satellite imagery, aerial imagery or other wide-area observation tools. In a 1540-relevant export and border-control context, remote sensing may be relevant to maritime routes, border regions, ports, storage facilities, industrial sites or logistical patterns. It can provide contextual information about movement and infrastructure, especially where physical inspection is difficult or where maritime and remote borders create practical monitoring challenges. Its value in this article is conceptual: routes are not only lines on a customs form, but spatial patterns that may interact with goods and data. Remote sensing can help illustrate how route awareness may be enriched by external evidence while still requiring careful interpretation.

The zero-shot capabilities of models such as VLMs and MLLMs may also be relevant. A model operating in a zero-shot setting may describe or classify something it has not been specifically trained to recognise as a controlled item, based on broader learned patterns. For export-control awareness, this is potentially relevant because ambiguity often concerns items that are not easily captured by a fixed list of examples. However, zero-shot outputs are also uncertain. They can be wrong, overconfident or sensitive to how the task is phrased. This reinforces the need for expert review.

An illustrative example is as follows: Consider a specialised industrial component with legitimate uses in ordinary manufacturing but also potential relevance in sensitive systems depending on its specifications and end-use. Traditional screening may see a generic product description and a broad commodity code. An AI-enabled system might compare the declared description with a product image, an X-ray scan, a technical datasheet and route information. The system might report that the visual features and internal structure are more complex than the description suggests, and that the route includes unusual trans-shipment. The system would output a structured interpretation for human review.

For non-technical audiences, the value of this approach lies in translating heterogeneous evidence into clear, reviewable questions: Does the image correspond with the declared description? Are the technical specifications consistent with the stated end-use? Does the route align with established trade patterns? Is relevant information missing, and do several weak signals become more significant when considered together? VLMs can compare visual and textual evidence, multimodal systems can integrate documents, images and sensor readings, and remote sensing can provide spatial and routing context. Together, these methods may support screening, prioritisation and interpretation, but their outputs should remain subject to assessment by appropriately trained officials, who retain responsibility for final legal and operational decisions.

Opportunities for Awareness, Monitoring and Implementation Support

The opportunities considered in this article relate to awareness, monitoring and implementation support rather than automated enforcement. AI-assisted document analysis could help organise the declarations, invoices, end-user statements, permits, licences, packing lists, technical datasheets and correspondence used in export- and border-control processes. Language models may support human review by summarising records, identifying missing information and highlighting inconsistencies across documents. Similar methods could also support education and training through simulated cases involving incomplete or conflicting documentation.

VLMs offer support by comparing visual and textual evidence. A VLM may assess whether a product photograph, X-ray scan or other image appears consistent with its declared description. Its role is not to determine whether an item is legally controlled, but to identify possible mismatches, uncertainty or visual features requiring expert attention. This distinction is important because visually similar items may have different technical specifications, while controlled and uncontrolled items may be difficult to distinguish from images alone. MLLMs could extend this approach by combining several forms of border-control evidence, including images, sensor readings, seals, manifests, vessel records, container histories and inspection notes. Organising these sources within a coherent analytical interface may help officials understand how different pieces of evidence relate to one another. However, multimodal analysis should remain a decision-support function: the relevance and reliability of each source must still be assessed by appropriately trained personnel.

AI-enabled analysis may also contribute to trade-route awareness. Routing and transaction data can be examined for patterns that differ from established trade activity, such as unexpected trans-shipment points, unfamiliar intermediaries or unusual combinations of goods and routes. Remote sensing may provide additional spatial context relating to ports, maritime movements, border regions or facilities. Such observations cannot, by themselves, establish that an activity presents a proliferation concern, because legitimate trade routes often change for commercial or logistical reasons. Their value lies in contextualising transactions and helping prioritise cases for further review.

Research-security education and interdisciplinary communication provide further opportunities. Researchers working in AI, engineering, chemistry, biology, materials science, robotics and remote sensing may contribute to UNSCR 1540-related discussions without becoming non-proliferation specialists. Accessible materials could explain how technical research connects with dual-use concerns, export-control ambiguity and responsible knowledge exchange. This is particularly relevant for early-career researchers who may develop advanced technical capabilities before encountering their wider policy and security implications.

Translation is required in both directions. Concepts such as VLMs, MLLMs may be unfamiliar to policy and practitioner communities. Conversely, technical researchers may have limited familiarity with related materials, means of delivery, strategic trade controls, trans-shipment and non-State actors. The goods–data–routes framework used in this article provides one way to connect these vocabularies by relating AI methods to practical sources of export- and border-control information.

These opportunities must be considered alongside technical and institutional limitations. AI systems can produce false positives, false negatives, unsupported explanations, biased outputs and poorly calibrated confidence scores. Sensor measurements may be noisy, remote-sensing imagery may be misinterpreted, and trade records may be incomplete. Performance may also decline when systems encounter rare items, unfamiliar languages, low-quality images or deliberate manipulation. AI outputs must therefore remain traceable to their supporting evidence and subject to human review.

Access to relevant technologies is also uneven. Advanced models, high-quality datasets, satellite imagery, specialised sensors and secure computing infrastructure may not be equally available across jurisdictions. A global approach should therefore avoid assuming that one technical configuration is suitable for every institutional or operational context. The emphasis should instead be placed on shared understanding, adaptable good practices and knowledge exchange.

A practical contribution could be the development of common educational scenarios. A simulated case might combine a broad HS code, an incomplete product description, a product image, an X-ray scan and a complex shipping route. Participants could examine what each source reveals, where uncertainty remains and how AI-generated assessments compare with expert judgement. Such scenarios would support awareness and interdisciplinary learning without disclosing sensitive operational information.

Illustrative AI-Assisted Workflows

The opportunities outlined in the previous section can be illustrated through simplified workflows that connect goods, data and routes with AI-assisted interpretation and human review. Table 1 summarises four examples. These workflows are intended to support awareness and education; they do not represent validated operational systems or recommendations for automated decision-making.

Table 1. Illustrative AI-assisted workflows for export- and border-control awareness and human review.

Inputs	Ambiguity	AI function	Limitations	Human responsibility
Declared description; product image; X-ray scan	The description is generic, incomplete or inconsistent with visible features.	A VLM compares visual and textual evidence and flags possible mismatches or inconsistencies.	Images may not reveal technical specifications, performance thresholds or legal status.	Review technical documentation, request clarification and determine whether further examination is justified.
Manifest; scan image; sensor readings; seal record; route history	Individual signals are incomplete, conflicting or inconclusive.	An MLLM organises the evidence and summarises relationships among the inputs.	Results depend on data quality, sensor calibration and model grounding; generated summaries may contain errors.	Assess the reliability of each source, verify highlighted inconsistencies and decide whether escalation is necessary.
Shipping records; port calls; vessel-tracking data; satellite or aerial imagery	An unusual route may reflect ordinary logistics or activity requiring further attention.	Data analysis and remote sensing provide spatial context and identify route deviations.	Coverage may be incomplete; imagery may be delayed or obscured; route anomalies do not establish a proliferation concern.	Compare the observed route with documentation, commercial context and other available information.
HS code; technical datasheet; product image; stated end-use	Commodity classification, technical characteristics and end-use information do not align clearly.	AI organises the case, identifies missing information and generates questions for review or training.	The model may overlook legal considerations, misinterpret technical terminology or express unwarranted certainty.	Apply authoritative guidance, explain relevant legal and technical considerations, and make the final assessment.

Across these examples, AI contributes by organising evidence, identifying possible inconsistencies and presenting questions for further review. Its output does not determine whether an item is controlled or whether a transaction presents a proliferation concern. Those assessments remain the responsibility of appropriately trained officials applying relevant legal, technical and contextual knowledge.

AI-Mediated CBRN Information Risks

AI may strengthen export and border controls by helping officials interpret incomplete, inconsistent or ambiguous information. At the same time, advanced AI models may lower barriers to accessing CBRN-related information. This dual effect requires careful assessment. Public discussion sometimes suggests that AI could enable individuals with limited expertise or infrastructure to develop CBRN capabilities. Such claims should be evaluated against the available evidence and the substantial material, technical and operational constraints involved.

Current evidence is better characterised as barrier lowering rather than capability replacement. AI systems can assist users in searching, summarising, translating, comparing and organising information from multiple sources. They may also support planning and troubleshooting. These functions could reduce some informational and organisational burdens, particularly where relevant knowledge is distributed across open literature, technical documentation and specialist practice. However, improved access to information does not replace the expertise, materials, equipment, facilities and practical skills required to translate information into physical capability.

The nature of these constraints differs across CBRN domains. Nuclear weapon development faces exceptionally high barriers, including access to fissile material, specialised facilities, advanced expertise and testing capabilities. AI-generated information does not remove these constraints. Radiological risks involve different material-security and public-safety concerns, but physical access to relevant materials remains central. In chemical and biological contexts, informational barrier lowering may be more plausible because aspects of scientific knowledge, synthesis and laboratory workflows are distributed across multiple sources. Nevertheless, a substantial gap remains between textual guidance and safe, reliable or harmful execution. Laboratory work depends on tacit skills, quality control, contamination management, safety procedures, materials handling and practical troubleshooting that cannot be reduced to written instructions alone.

Evaluations conducted by OpenAI and RAND illustrate this cautious approach (Mouton et al., 2023; OpenAI, 2024a). OpenAI examined whether access to an AI model improved performance on biological-threat-related tasks compared with internet-only resources, while emphasising that information forms only one part of a broader capability chain involving expertise, materials, equipment and execution (OpenAI, 2024a). RAND similarly assessed whether large language models could assist biological-attack planning, identifying potential informational support without concluding that these systems provide operational capability (Mouton et al., 2023). Together, these studies support proportionate awareness of possible barrier-lowering effects while maintaining a clear distinction between information access and real-world capability.

Frontier AI developers have also incorporated CBRN misuse into their safety frameworks. OpenAI's preparedness framework includes CBRN-related risk categories and capability thresholds (OpenAI, 2023), while Anthropic has described strengthened safeguards for models where biological or broader CBRN misuse risks cannot be confidently excluded (Anthropic, 2025). These developments indicate that the relationship between AI capabilities, safety evaluation and WMD-related misuse is becoming an established concern within technical governance. They do not, however, demonstrate that AI directly provides non-State actors with the capability to develop or produce such weapons.

The relevance to export and border control lies in the interaction between information and physical access. AI-assisted information processing may support design research, procurement analysis and the identification of materials, components or services relevant to sensitive activities. The same capabilities may also help legitimate organisations assess transactions, improve documentation, identify inconsistencies and train compliance personnel. AI can therefore create both risks and implementation-support opportunities.

Within the article's goods–data–routes framework, AI-mediated information represents only one part of the wider risk environment. Information may influence how goods are identified, acquired or moved, but access to related materials, equipment, technologies and services remains subject to supply-chain conditions and national controls. Appropriate and effective export and border controls therefore remain central to limiting the translation of AI-assisted information into practical CBRN capabilities.

Responsible Use, Uncertainty and Human Judgement

AI-enabled systems may help make export- and border-control ambiguity more visible, but they do not replace expert decisions, which can have legal, commercial, diplomatic and security consequences. Human review, institutional accountability and domain expertise remain essential. One reason is uncertainty. AI systems may generate plausible explanations that are not grounded in reliable evidence. For example, a VLM may misidentify a component, and a language model may misunderstand a technical term. A remote-sensing interpretation may confuse ordinary industrial activity with suspicious activity. These errors can matter. Responsible awareness therefore includes understanding model limits. False positives and false negatives have different implications. A false positive may burden legitimate trade, create unnecessary scrutiny or reduce trust in technology. A false negative may miss a case that deserves attention. In low-prevalence domains, even a model with high apparent accuracy can generate many false alarms. This statistical reality is often unfamiliar to non-technical audiences and should be part of education. AI outputs need to be interpreted with base rates, context and uncertainty in mind.

Explainability is also important. In export and border-control contexts, it is rarely enough for a system to say that something is high risk. A human reviewer needs to understand why. Did the system notice a mismatch between image and description? Did it identify missing technical specifications? Did it compare the route with ordinary trade patterns? Did it associate a term with a control-list concept? Explanation does not guarantee correctness, but it allows the reviewer to question, validate or reject the output.

Data quality is a further constraint. Customs and trade data may contain errors, abbreviations, inconsistent languages, legacy formats or missing fields. Sensor data may be affected by calibration, environmental conditions or operator practice. Images may be low resolution or occluded. Remote-sensing data may be delayed or cloud-covered. AI systems trained or evaluated on clean benchmark data may behave differently in messy operational contexts. Awareness-building should therefore include data realism.

Adversarial behaviour must also be considered. Non-State actors or proliferation networks may misdescribe goods, split orders, use front companies, alter routes, disguise end-users or manipulate images and documents. AI systems can be part of the response, but they can also be targeted. Models may be probed, fooled or used by malicious actors to improve deception. This dual-use dynamic reinforces the need for layered interpretation rather than reliance on a single tool.

Human-in-the-loop language can be too vague unless it specifies the role of the human. In this article, the relevant human may be a customs officer, export-control officer, technical expert, compliance analyst, laboratory safety specialist, maritime analyst or policy practitioner. Their task is not merely to approve an AI output. It is to interpret evidence, apply legal and technical knowledge, request further information, escalate uncertain cases and communicate decisions through appropriate institutional channels. AI support has value only if it fits these human workflows.

Responsible use also requires sensitivity to rights, privacy, commercial confidentiality and legitimate scientific exchange. Export and border-control awareness sits at the intersection of security and trade. Remote sensing, data analytics and AI screening can raise concerns about surveillance, data protection or commercial sensitivity. These concerns should be acknowledged, and they reflect the institutional reality that UNSCR 1540 is a legal and political framework, while AI systems are technical artefacts requiring careful integration into human processes.

Education, Research Engagement and Expert Networks

In this article, we identify where AI research may be relevant to export and border-control discussions under UNSCR 1540, with an emphasis on academic perspectives and awareness-building. Education, research engagement and expert networks are also important in this context. The target audience includes not only 1540 specialists, but also researchers in computer science, engineering, life sciences, remote sensing, law, policy and international security who may encounter dual-use questions in their work. Education is important, since many researchers do not initially see the connection between their technical work and non-proliferation. For example, a computer scientist working on VLM may not think about customs declarations; a remote-sensing researcher may not think about strategic trade routes; a law or policy researcher may not know how AI systems interpret images. The goods–data–routes organising framework provides an accessible entry point for these researchers.

Research engagement can take several non-prescriptive forms. Universities and research networks may host workshops that introduce UNSCR 1540 concepts to technical audiences. Technical groups may develop neutral explainers on AI, sensing or remote monitoring for policy audiences. Interdisciplinary teams may create simulated case studies for training. Early-career researchers may participate in writing groups, short courses or policy-facing briefings. Short courses on dual-use awareness, research security, export-control concepts and AI limitations could help researchers communicate across communities.

Expert networks are particularly useful because no single discipline owns the problem. For example, export and border-control ambiguity involves law, customs practice, data science, engineering, chemistry, biology, logistics, finance, security studies and international institutions. A technical model may be elegant but irrelevant if it does not match practitioner needs. A policy concept may be sound but technically unrealistic if it misunderstands data or model limits. Networks can help identify these mismatches early. Industry is also part of the knowledge environment. Logistics companies, manufacturers, cloud laboratories, additive manufacturing services, AI developers, satellite data providers and compliance software companies all hold relevant expertise, and knowledge exchange across academia, industry, international organisations and practitioners may improve understanding of emerging risks and support tools. This is especially true where new service models complicate end-use awareness.

By using the goods–data–routes framework, this article translates selected AI research into UNSCR 1540-relevant language and connects export-control ambiguity with technical interpretation tools. It thereby provides a starting point for education, research engagement and expert networks concerned with UNSCR 1540 in an evolving technological landscape.

Conclusion

UNSCR 1540 provides an enduring and technology-neutral framework for preventing non-State actors from acquiring nuclear, chemical and biological weapons, related materials and means of delivery. Emerging technologies do not alter this core purpose, but they are changing the practical environment in which export and border controls operate. Digital transfers, service-based production, complex supply chains and increasingly information-rich trade processes make it more difficult to assess goods, transactions and patterns of movement through control lists or individual data sources alone. Using the goods–data–routes framework, this article has shown that practical ambiguity often arises when physical items, supporting information and trade routes do not align clearly. Control lists remain essential for establishing legal requirements and identifying known sensitive items. However, effective assessment may also depend on technical specifications, product images, documents,

sensor readings, end users, service relationships and route context. Ambiguity cannot be eliminated through automated analysis, but it may be made more visible and manageable through the structured interpretation of these different sources.

AI-enabled systems may support this process in several ways. VLMs can compare images with textual descriptions; MLLMs can organise evidence drawn from documents, images, sensor data and structured records; and remote-sensing methods can provide additional spatial and route context. These capabilities may help identify inconsistencies, connect weak signals and prioritise cases for further examination. They do not determine whether an item is legally controlled, establish proliferation intent or replace the authority and judgement of export-control and border-control professionals. Their appropriate role is therefore to support, rather than automate, assessment and decision-making.

The dual-use character of AI also requires careful attention. AI may lower some informational, planning and organisational barriers relevant to CBRN misuse, particularly in chemical and biological contexts. Nevertheless, access to information remains distinct from the ability to develop a practical capability, which continues to depend on materials, equipment, facilities, tacit expertise, safety conditions, testing and supply-chain access. This distinction supports proportionate risk communication while reinforcing the continuing importance of appropriate and effective controls over related materials, technologies and services.

The main contribution of this article is to translate relevant AI research into language that can support UNSCR 1540 awareness, education, research engagement and implementation-support discussions. Rather than presenting AI as a technological solution to proliferation, the article identifies areas in which it may help human experts interpret increasingly complex evidence. This provides a basis for further interdisciplinary dialogue among technical researchers, practitioners, industry, international organisations and the wider UNSCR 1540 community. The central conclusion is that AI cannot remove ambiguity or replace expert judgement, but it may help ensure that important relationships among goods, data and routes are more systematically recognised and examined.

Appendix

Vision-Language Models

A vision-language model (VLM) learns relationships between visual inputs and natural language. The term covers several architectures, ranging from discriminative models that compare images and text to generative systems that produce captions or answer questions about images. A typical VLM contains a visual encoder, a text encoder or language model, and an alignment or fusion mechanism that connects the two representations.

CLIP is a prominent example of a discriminative VLM (Radford et al., 2021). It uses contrastive learning: matching image-caption pairs are moved closer together in a shared embedding space, while non-matching pairs are moved farther apart. Once trained, the model can retrieve images from text, retrieve text from images, or classify an image using category names supplied as natural-language prompts. This enables zero-shot transfer, in which a model is applied to categories that were not explicitly represented as fixed labels during task-specific training.

An embedding is a numerical vector that represents selected semantic properties of an input. Image-text consistency can therefore be estimated by comparing the image and text embeddings, commonly using cosine similarity. A high similarity score suggests that the two inputs are semantically aligned while a low score indicates possible mismatch, and an example can be found in Fig. 2. The score is probabilistic and task-dependent: it does not prove that either input is correct or false, and performance can decline when images, terminology or concepts differ from the training data.

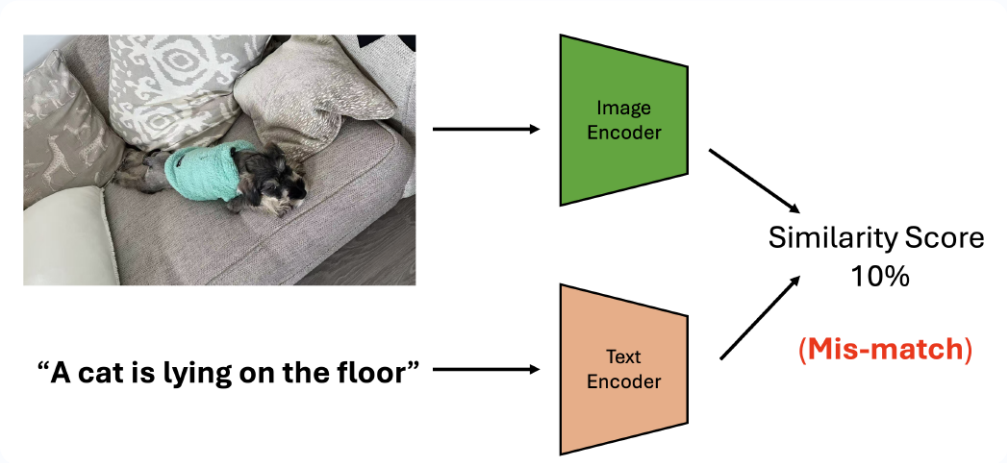

Figure 2. Image-text similarity calculation for consistency checking using a vision-language model, i.e., CLIP (Radford et al., 2021).

Generative VLMs connect visual representations to a language-generation component. BLIP2, for example, links a pretrained image encoder and a pretrained Large Language Model (LLM) through a lightweight Querying Transformer (Q-Former) that learns to bridge the two modalities (Li et al., 2023). Visual question answering (VQA) asks a model to generate an answer conditioned on both an image and a natural-language question. Formally, the system estimates an answer from the conditional relationship between image and question, rather than from either input alone. VQA may be closed-ended, with answers selected from a predefined set, or open-ended, with answers generated as free text. Reliable performance requires visual grounding: the model must connect words in the question to relevant regions, objects or relationships in the image, and an example of VQA can be found in Fig. 3.

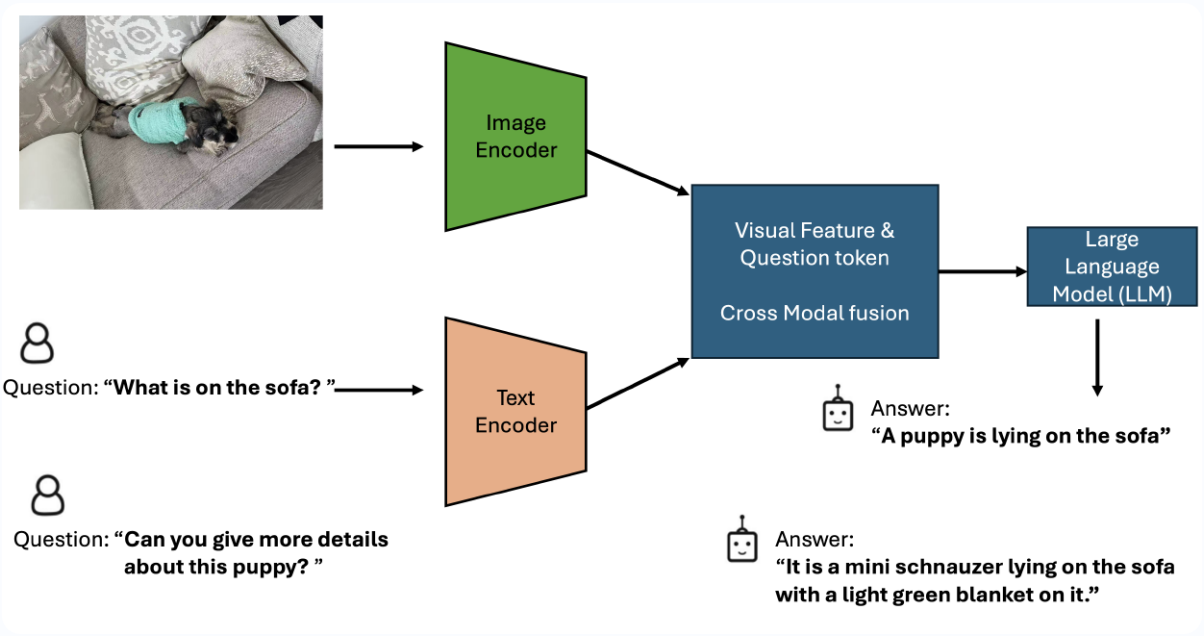

Figure 3. a visual question-answering example using a vision-language model, i.e., BLIP-2 (Li et al., 2023).

Multimodal Large Language Models

An MLLM uses a large language model as a general-purpose reasoning and generation component while accepting one or more non-text modalities (OpenAI, 2024b). A generative image-language system may therefore be both a VLM and an MLLM. By contrast, a dual-encoder model such as CLIP is normally described as a VLM but not an MLLM because it does not use a generative large language model as its central interface.

MLLMs commonly use modality-specific encoders to transform heterogeneous inputs into numerical representations. A projector, adapter or cross-attention module then maps those representations into a form that the language model can process. Depending on the architecture, multimodal information may be inserted as tokens, accessed through cross-attention, or combined through an external tool. Training often includes three stages: pretraining the individual components, aligning non-text representations with language, and instruction tuning on examples of multimodal tasks.

Possible modalities extend beyond image and text. ImageBind demonstrates a shared representation space for images, text, audio, depth, thermal measurements and inertial measurement unit data (Girdhar et al., 2023). PaLM-E incorporates visual observations and continuous sensor-state information into a language-model interface (Driess et al., 2023). Tabular data can be supplied by serialising rows and fields into text, as demonstrated by TabLLM, or by using a dedicated table encoder and projection layer (Hegselmann et al., 2023). Fig. 4 provides an example of MLLM combining image text, sensor, audio and tabular inputs with various downstream tasks such as question answering.

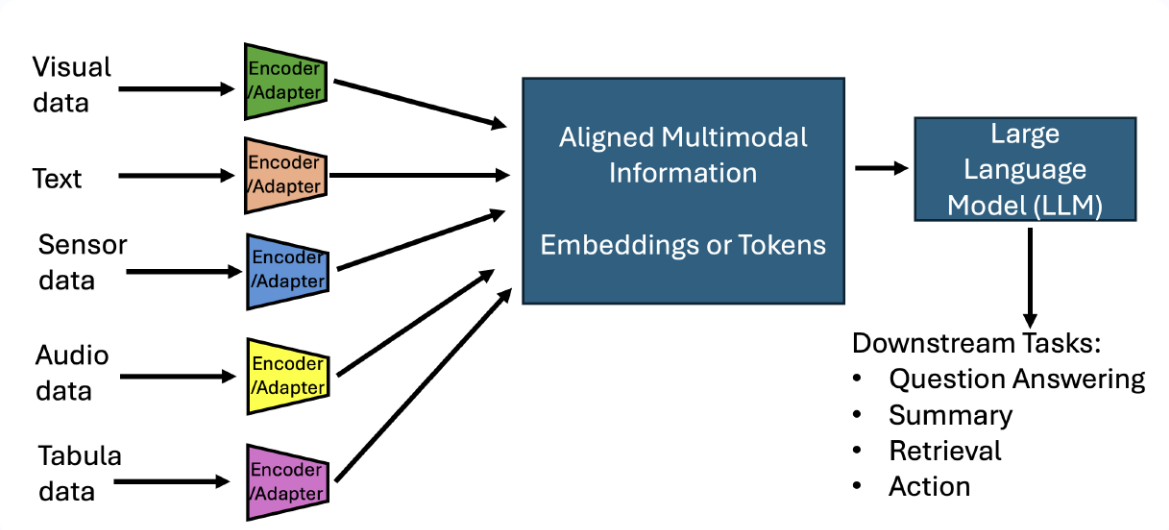

Figure 4. Example architecture of a MLLM accepting image, text, sensor, audio and tabular inputs.

Multimodal fusion may occur at different stages. Early fusion combines low-level inputs or tokens before substantial modality-specific processing. Intermediate fusion combines learned features, often through cross-attention or a shared latent space. Late fusion combines independently generated predictions or scores. No strategy is universally best: the appropriate design depends on whether modalities are synchronised, whether some inputs are missing, and whether the task requires fine-grained interaction between them.

The composition of multimodal training data strongly influences what a model learns. Web-scale image-text collections provide breadth but may contain noisy captions, duplicated material, cultural bias or weak correspondence between an image and its accompanying text. Audio, sensor and tabular datasets are often smaller

and collected under different sampling conditions, creating imbalance among modalities. Furthermore, an alignment objective may place related inputs close together without preserving every detail required by a downstream task. Broad pretraining should therefore be followed by task-specific evaluation and, where appropriate, carefully controlled adaptation.

MLLM outputs remain probabilistic. Important limitations include hallucination, weak calibration, spurious correlations, modality imbalance, missing or noisy inputs, and domain shift. A model may rely on language priors instead of visual evidence, ignore a weak sensor stream, or generate a fluent explanation unsupported by the inputs. Evaluation should therefore test each modality separately and in combination, include cases with conflicting or missing evidence, and distinguish linguistic plausibility from factual and perceptual accuracy. The central advantage of an MLLM is its ability to provide a common language-based interface across heterogeneous information; its central challenge is ensuring that generated outputs remain grounded in the supplied evidence.

References

Anthropic. (2025, May 22). Activating AI Safety Level 3 protections. https://www.anthropic.com/news/activating-asl3-protections

Driess, D., et al. (2023). PaLM-E: An embodied multimodal language model. Proceedings of the 40th International Conference on Machine Learning.

Girdhar, R., et al. (2023). ImageBind: One embedding space to bind them all. Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition.

Hegselmann, S., et al. (2023). TabLLM: Few-shot classification of tabular data with large language models. Proceedings of the 26th International Conference on Artificial Intelligence and Statistics.

Li, J., et al. (2023). BLIP-2: Bootstrapping language-image pre-training with frozen image encoders and large language models. Proceedings of the 40th International Conference on Machine Learning.

Mouton, C. A., Lucas, C., & Guest, E. (2023). The operational risks of AI in large-scale biological attacks: A red-team approach (Research Report RR-A2977-1). RAND Corporation. https://www.rand.org/pubs/research_reports/RRA2977-1.html

OpenAI. (2023, December 18). Preparedness framework (Beta). https://cdn.openai.com/openai-preparedness-framework-beta.pdf

OpenAI. (2024a, January 31). Building an early warning system for LLM-aided biological threat creation. https://openai.com/index/building-an-early-warning-system-for-llm-aided-biological-threat-creation/

OpenAI. (2024b, May 13). Hello GPT-4o. https://openai.com/index/hello-gpt-4o/

Radford, A., et al. (2021). Learning transferable visual models from natural language supervision. Proceedings of the 38th International Conference on Machine Learning.

Stockholm International Peace Research Institute. (2002). Multilateral export controls. In SIPRI yearbook 2002: Armaments, disarmament and international security. https://www.sipri.org/yearbook/2002/14

United Nations Interregional Crime and Justice Research Institute. (2024). UNSCR 1540: Operative paragraph 3. https://unicri.org/sites/default/files/2024-09/UNSCR1540-Operative-Paragraph-3.pdf

United Nations Security Council. (2004). Resolution 1540 (2004), S/RES/1540 (2004). https://docs.un.org/en/S/RES/1540(2004)

United Nations Security Council 1540 Committee. (n.d.). Voluntary technical reference guide: Export controls. https://www.un.org/en/sc/1540/national-implementation/vTRG-ExpCtrl-1.shtml

Author Bios

Yu Guan is an Associate Professor in the Department of Computer Science at the University of Warwick and leads the Ubiquitous and Visual Computing Lab. His research focuses on machine learning, computer vision and multimodal sensing, including cross-modal learning and the interpretation of image, video and sensor data. He is an associate editor of ACM on Interactive, Mobile, Wearable and Ubiquitous Technologies (IMWUT).

Sai Gu is Professor of Chemical Engineering and Deputy Pro-Vice-Chancellor at the University of Warwick. He is a Fellow of the Royal Academy of Engineering (FREng), and his research centers on sustainable manufacturing, clean energy, and process digitalization. He is the leader of UNSCR 1540 TechHorizon Initiative, aiming at connecting emerging technology research with the UNSCR 1540 community, improving understanding of how technological change may create new proliferation risks while also supporting awareness, monitoring, education and implementation-related discussions. He has extensive experience in leading multidisciplinary and international research collaborations, providing a broad engineering and institutional perspective for this TechHorizon Initiative.

ARTICLE 05

Aerial Drones and "Means of Delivery" in the Context of UN Security Council Resolution 1540 (2004)

Matthew J. F. Healy
Associate Professor of Nuclear, Radiological and Non-Conventional IED Defence
Centre for Defence Engineering and Physical Science, Cranfield University

Andrew Horton
Visiting Senior Research Fellow
King's Collage London, Centre for Science & Security Studies (CSSS), Department of War Studies



Abstract

Aerial drones have become a highly accessible technology used for recreation, public service and commercial applications.[1] They may be equipped with a wide range of sensors, communication systems, carry payloads of different types, and can even carry dispersion systems for precision agricultural purposes. Many can carry significant payloads, at relatively low cost which in some applications can make them disposable. They are also difficult to detect, stop and to trace to their operators. This paper considers the physical capacity of aerial drones and the typical planning considerations of malicious or violent non-state actors, in terms of their potential use for the dispersal of chemical, biological and radiological agents. The article also considers possible existing international legal instruments and normative arrangements that could be used as a regulatory framework aimed at preventing the use of aerial drones as a delivery mechanism for weapons of mass destruction, especially in the context of meeting state obligations under United Nations Security Council resolution 1540 (2004) in regard to WMD means of delivery [1].

1. There are a wide range of terms used for such "drones" such as Unmanned Aerial Vehicles or Un-crewed Aerial Vehicles (UAVS), Unmanned Aerial Systems (UAS), Remotely Piloted Aerial System (RPAS), the commonly term "drone" is used to cover all types of vehicle or system, be they remotely piloted or operated semi or fully autonomously.

Introduction

Aerial drones have in recent times changed the nature of many aspects of security and defence [2]. Their uses range from simple incursions of privacy as aerial platforms for obtaining video [3], to platforms that inexpensively carry explosive payloads for lethal purposes by drug gangs [4,5] and are used in numbers that overwhelm state level defences [6,7]. The concept, and capacity for remotely controlled airborne delivery of lethal effects is far from new; small remote controlled combustion engine helicopter designs have enabled this capacity for over 60 years, an example of this is the US Navy's QH-50 Drone Anti-Submarine Helicopter (DASH) designed in the 1950s and where the lethal effect was via a Mk. 57 Nuclear Depth Charge (which had a maximum yield of 20 kt). However, many of the issues inherent in these early drones have been overcome with more modern designs through developments in technology. The use of electric motors and high-capacity batteries, geospatial tracking, auto route control, stabilisation and first-person view (FPV) navigation have made the use of drones easier, their costs lower and far more accessible, and attractive to non-state actors. The ever greater range of uses of remotely piloted, semi and fully autonomous aerial drones envisaged as part of the so-called Advanced Air Mobility (AAM) concept, and in particular Urban Air Mobility (UAM), will make it potentially far easier for a range of malicious or violent non-state actors (M/VNSAs) to get hold of advanced drone technologies and access capabilities that up until only recently were the preserve of the state.

The development in the availability and use of drones has arguably outpaced the corresponding development of legal control and regulations around their use, especially in the context of security and the potential use by those with criminal or malicious intentions. When it comes to the use of drones as a means of delivery of the full range of chemical, biological, radiological and nuclear (CBRN) agents and weapons then there is a complex, but incomplete, range of international measures both in terms of "hard" laws and "soft" political arrangements, around which national legislative and regulatory measures can be built. This has implications for states in regard to meeting their obligations under UN Security Council resolution 1540 (2004) in relation to addressing means of delivery for nuclear, chemical or biological weapons. The resolution requires that states undertake measures in regard to the potential proliferation, by and to non-state actors, of "*missiles, rockets and other unmanned systems capable of delivering nuclear, chemical, or biological weapons, that are specially designed for such use*". Whilst the resolution limits its specific requirement through the use of the term "*specially designed for such use*", this should be seen as a minimum requirement, and that there is a greater risk when it comes to non-state actors around the potential modification of commercially available drones and leisure drones as an improvised means of delivery. Therefore, to fully and effectively prevent non-state actors from acquiring aerial drones as a WMD means of delivery, then these risks should also be considered when it comes to developing possible legal and regulatory controls.

Malicious / Violent Non-state Actors (M/VNSAs) Methodology

Acts perpetrated by a range of M/VNSAs may occur for many reasons, ranging from inappropriately protesting, achieving financial gain, personal revenge to terrorism. There are many phases to such acts where drones can support actors operating outside of state control to achieve serious widespread harm.

Although the motivations for the actions of M/VNSA are broad; the field of terrorism provides a mature and useful broad-based framework within which the phases and criteria of most other non-state acts fit [8, 9]. The terrorist attack planning cycle features elements ranging from reconnaissance and target selection through to delivery, release, evasion and exploitation.

For the purpose of encouraging the prevention of aerial drone technology being abused by non-state actors it is useful to expose the leverage that a delivery system achieves for hazardous materials, where delivery and material combine to form a weapon of mass destruction. The following section shows how both the motivation of a M/VNSA and the harmful physical consequences may be increased through phases of the attack cycle. The capacity for small drones to carry surveillance equipment and conduct target selection and reconnaissance is well known [10,11], often considered, and are generic to many attack types so they will not be the focus.

M/VNSA self-preservation and spatially constraining CBRN effects

One facet of M/VNSA methodology with any kind of weaponry is that actors do generally factor in their own safety. Suicide actors exist, but they still wish to preserve themselves long enough to achieve their aims and recognize success. Methods for weaponizing CBRN materials, use and their effects are not well understood nor are they easy to control, and this contributes a high degree of uncertainty into a protagonist's decision making. The most obvious phase where drones can separate the protagonist from harm is delivery to target. CBRN materials do not always need to be positioned at the target, instead they were traditionally released some distance away from the target to be carried by wind as occurred in World War 1. However, WW1 showed how a changing wind direction or drop in wind speed could harm the protagonists themselves [12]. More recent use of chlorine by the Liberation Tigers of Tamil Eelam in Sri Lanka 1990 showed how harm that was intended for the Sri Lankan army also fell on local civilians. This severely undermined the support base and legitimacy of the LTTE non-state actor group and ended their use of chlorine [13]. Drones in warfare are being used to linger near potential targets ready to seize attack opportunities. Drones are currently being used in Ukraine to precisely target individuals with kinetic effects, and without reference to drones, both governments are claiming the other has used chemical weapons, with denials [14]. However, considering that UNSCR1540 applies to non-state actors then it is worth noting that a state's soldier who operates outside of his authority is still usually considered a state actor (as are combatants originating from another state) according to the Articles on Responsibility of States for Internationally Wrongful Acts adopted by the International Law Commission in 2001 [15]. Whilst any acts involving chemical weapons by either side in Ukraine are outside of the scope of UNSCR1540, the concepts, experiences, and especially the rapidly evolving drone technology developed in this conflict is vulnerable to reaching non-sate actors through the scale of the conflict and the range of state, non-state and industrial relationships that have developed because of it. The need for drones to conduct sampling in hazardous environments is driving technology forward with significant funding from governments [16]. Simultaneous to achieving agent delivery, the methods and technology used for surveillance could be highly effective in exploitation by streaming a terror event live to the public.

Access

One of the clearest advantages of drones for the MVNSA is that they are excellent at circumventing traditional perimeter security. However, the way they can achieve this couples well with exploiting harmful characteristics of CBR agents. An incident in Tokyo in 2015 highlighted properties of drones relevant to the use of CBRN weaponry. A drone was landed on the roof of an official residence of Prime Minister Shinzo Abe. The site is one that would be difficult to access by an unauthorised human. The drone was landed at 3:30am on April 9th but was not removed until April 22nd. The drone's cargo was radioactive waste material obtained from the Fukushima region [17]. An analogous exposure event (not involving drones) occurred in China [18, 19] where a nuclear expert exposed a colleague to approximately 1 Sv of radiation by concealing Iridium-192 in his office ceiling for over 2 months. The deliberate attack also injured 73 other hospital staff and patients including a woman who reportedly nearly miscarried. Although this incident involved higher activity material than the Tokyo drone incursion, it indicates the potential mass effect of an attack from a radioactive source.

Another facet of the 2015 Tokyo incursion is that the protagonist abandoned the drone; unlike humans modern drones can be cheap and available enough to be expendable. This has significant impact on potential attack methodologies because the exit phase (following delivering an agent to its point of release) becomes optional. The choices of entry and exit routes for a mal-actor are not independent.

Brute Force Access and Speed

Drones are difficult to stop from accessing protected spaces, even if they are expected or identified. Drones have extensively disrupted airport activity at Gatwick airport in 2018 [20], also a drone overflew crowds at the British King's coronation in 2024 [21] and they are commonly used to deliver contraband (particularly drugs) into prisons [22,23]. Drone countermeasures exist, but they are often highly localised There is a rapid evolution of technology and tactics from both protagonists and protective forces in relation to drone deliveries in prisons.

Drones have been used to access sensitive areas in a wide range of events involving high concentrations of people. There have been at least ten significant incursions of drones over sporting stadia in the last six years, including fixtures at high volume international stadia [24]. The mitigation of these incursions (if at all) has usually been the suspension of the sporting activity rather than any measure to protect the crowd. A similar number of instances have occurred at political events and festivals with an event in Brighton involving a drone flying particularly low over a crowd of 8500 people before reaching the performance stage [25]. In ground based conventional marauding attack events the public have a chance to respond and defeat an attacker [26] whereas few options exist for the public to defeat an overhead attack.

Payload

It is sometimes thought that aerial drones cannot deliver significant quantities of harmful material unless that material were very highly toxic. Some of that perception arises from comparisons with the very large volumes of toxic industrial material that have been involved in accidents like the Methyl Isocyanate release in Bhopal 1984 [27] or that are regularly used in traditional ground-based tractor pulled spraying of crops which carry volumes of liquid 2 to 3 orders of magnitude more than typical airborne drones can carry. The high concentration of people at sporting stadia, festivals, or political rallies or protests mean that aerial drones need only cover much smaller areas to harm people than to spray crops. Discussing the efficacy of specific poisons against humans is beyond the scope of this paper, but the volume of liquid that may be required by a drone to uniformly cover 5 acres typical of humans at stadia would be 100 to 2000 litres. Whilst most of this payload range is still beyond most drones it is achievable.

CBRN Means of Delivery under International and Domestic Law

Whilst resolution 1540 (2004) does not provide the specific details of how the measures that states are required to implement are implemented, it does require that those measures that have been adopted by states through international "hard" law instruments such as treaties, and "soft" normative arrangements can be applied to non-state actors through domestic legislation. The resolution is very much seen as being complementary, and not to replace or undermine those existing instruments and arrangements. When it comes to the question of CBRN means of delivery, whilst there is currently no specific aerial drone treaty or arrangement, there are nevertheless a number of instruments and arrangements that can be helpful when looking to develop a legal framework.

Hard International Law

When it comes to the question of legally binding international treaties and conventions, then there are several in the area of arms control and terrorism that are helpful as they either make a specific reference to “means of delivery”, or at least there is potential to apply them through interpretation of terms such as “devices” or “equipment”, and in the case of aviation law the definition of “aircraft”.

1972 Biological and Toxin Weapons Convention (BTWC) [28]

Article 1 of the BTWC has the following undertaking which is placed on state parties to the convention -

Each State Party to this Convention undertakes never in any circumstances to develop, produce, stockpile or otherwise acquire or retain:

(1) microbial or other biological agents, or toxins whatever their origin or method of production, of types and in quantities that have no justification for prophylactic, protective or other peaceful purposes;

(2) weapons, equipment or means of delivery designed to use such agents or toxins for hostile purposes or in armed conflict.

1997 Chemical Weapons Convention (CWC) [29]

The CWC in Article 2 has the following definition -

"Chemical Weapons" means the following, together or separately:

1. *Toxic chemicals and their precursors, except where intended for purposes not prohibited under this Convention, as long as the types and quantities are consistent with such purposes;*

2. *Munitions and devices, specifically designed to cause death or other harm through the toxic properties of those toxic chemicals specified in subparagraph (a), which would be released as a result of the employment of such munitions and devices;*

3. *Any equipment specifically designed for use directly in connection with the employment of munitions and devices specified in subparagraph (b).*

1997 International Convention for the Suppression of Terrorist Bombing [30]

Article 1 Paragraph 3(b) of the convention talks about *"weapon or device that is designed, or has the capability, to cause death, or serious injury or substantial material damage through the release, dissemination or impact of toxic chemicals, biological agents or toxins or similar substances or radiation or radioactive material."*

2010 International Convention on the Suppression of Unlawful Acts Relating to International Civil Aviation (Beijing Convention) [31]

Article 1(1)(g) requires that states party to the convention make it a criminal offence for a *"... person unlawfully and intentionally ... releases or discharges from an aircraft in service any BCN weapon or explosive, radioactive, or similar substances in a manner that causes or is likely to cause death, serious bodily injury or serious damage to property or the environment".*

It is perhaps the 2010 Beijing Convention, which is focused on aviation law rather than arms controls or counter-terrorism, that is the most relevant to the issue of the use of drones as a means of delivery in the context of resolution 1540 (2004), as the definition of "aircraft" can include "drones". The foundational document of aviation law, including the Beijing Convention, is the *1944 Convention on International Civil Aviation* (Chicago Convention) [32], and Article 8 of that convention, concerning *pilotless aircraft*, states that *"no aircraft capable of being flown without a pilot shall be flown without a pilot over the territory of a contracting State without special authorization by that State and in accordance with the terms of such authorization".* This clearly shows that the term "aircraft" as used in international civil aviation law, whilst more often viewed as being concerned with manned aircraft, can also be apply to aerial drones.

Soft International Political Agreements

In the case of normative arrangements, then the focus shifts more towards the area of non-proliferation and in particular the export control of means of delivery. There are three Multilateral Export Control Regimes (MECRs) that control the transfer of aerial drones and related dual-use technology, these are, the *Missile Technology Control Regime (MTCR)*, the *Wassenaar Arrangement (WA)* and the *Australia Group (AG)*, each of these are described in more detail below:

Missile Technology Control Regime [33]

The MTCR is the only MECR that is focused solely on the issue of WMD means of delivery, and controls aerial drones in the following manner - Complete unmanned aerial vehicle systems (including cruise missiles, target drones and reconnaissance drones) that are capable of delivering at least a 500 kg payload to a range of at least 300 km, these are known as *Category-1* and are the most restricted type of systems. For those capable of a range equal to or greater than 300 km, but below the 500kg payload of Category-1, these fall into *Category-2* and are subject to less stringent requirements around their transfer. In addition, there are controls on unmanned aerial vehicle systems, not specified in either category-1 or 2, but have either an autonomous flight control and navigation capability, or capable of controlled flight out of the direct vision range of a human operator; and "*incorporating an aerosol dispensing system/mechanism with a capacity greater than 20 litres, or designed or modified to incorporate an aerosol dispensing system/mechanism with a capacity greater than 20 litres.*" The accompanying technical note notes that -

> This includes the possibility of aerosol injection into the combustion exhaust vapour and into the propeller slip stream. Where an aerosol consists of particulate or liquids other than fuel components, by-products or additives, as part of the payload to be dispersed in the atmosphere. Examples of aerosols include pesticides for crop dusting and dry chemicals for cloud seeding.

Wassenaar Arrangement [34]

The WA is concerned with the transfer of both military and related dual-use systems. In the case of non-military drones then the regime controls - *UAVs or unmanned airships, designed to have controlled flight out of the direct 'natural vision' of the 'operator' and having any of the following: a maximum endurance greater than or equal to 30 minutes but less than 1 hour and designed to take-off and have stable controlled flight in wind gusts equal to or exceeding 46.3km/h (25 knots), or a maximum endurance of 1 hour or greater.*

Australia Group [35]

> The focus of the AG is to control the transfer of items concerned with chemical or biological weapons and it is in this context that the regime controls "spraying or fogging systems and components" in particular the AG controls - *"complete spraying or fogging systems, specially designed or modified for fitting to aircraft, lighter than air vehicles or UAVs, capable of delivering, from a liquid suspension, an initial droplet VMD of less than 50 microns at a flow rate of greater than two litres per minute"* as well as, *"spray booms or arrays of aerosol generating units, specially designed or modified for fitting to aircraft, lighter than air vehicles or UAVs".* The accompanying technical note states that it does not control those systems that are *"demonstrated not to be capable of delivering biological agents in the form of infectious aerosols",* which indicates that the intent of the control is about the delivery of biological agents.

The usefulness of the MECRs, beyond that of export controls, is that they provide more technical details about the type of drone systems that are potentially useful for MVNSAs, and these technical details can be incorporated into domestic legislation or regulations, to help narrow the scope of any controls.

Domestic Regulation

One example of domestic legislation relevant to resolution 1540 (2004), which links drones and the delivery of chemical or biological agents, and radioactive substances, is that of Singapore's Air Navigation Act 1966, which under Sections 34 and 35, prohibit the carriage of dangerous materials on unmanned aircraft, which includes "any biological agent, biological agent waste or toxin within the meaning of the Biological Agents and Toxins Act 2005" and "any radioactive material, radioactive substance or radioactive waste within the meaning of the Radiation Protection Act 2007", then section 34 concludes with "any other hazardous material (whether gaseous, liquid or solid) that is prescribed, in a notification in the Gazette, by the Minister charged with the responsibility for homefront security to be a prohibited item for the purposes of this section". Section 35 furthermore covers the "discharge" of such prohibited materials from unmanned aircraft.

The Growing Need to Develop Regulations

One commercial area that has exploited the use of civilian aerial drones over the last few decades is precision agriculture. The sector has used drones for a number of applications, such as the Yamaha R-Max unmanned helicopter which is often employed for aerial spraying. The R-Max was developed in the 1990s in response to the needs of Japan's small farm based rural agriculture, the R-Max has now been approved for commercial operations in both the US and New Zealand. [36,37] The R-Max was a development of the Yamaha R-50 Remote Control Aerial Spraying System (RCASS) which started operations in 1987, and had a payload capability of 20kgs [38], (it's no coincidence when compared to the 20 litres as adopted by MTCR for its UAV controls as mentioned above), this compares with the R-Max's maximum payload of around 30kgs.

Going beyond precision agriculture, over recent years there has been increasing interest in the potential development of so-called Advanced Air Mobility (AAM), which includes Urban Air Mobility (UAM), which is the use of autonomous or remotely controlled drones to undertake a number of roles within the urban environment, such as the delivery of packages, logistic (first-mile/last-mile) operations, and even the transport of people in passenger drones. The development of the commercial use of aerial drones has presented a challenge to regulators, given the ever-expanding commercial applications and leisure use of aerial drones this regulatory challenge is only likely to grow over the coming years. Whilst most of the focus of the regulators has been on the safety aspects of using aerial drones in commercial operations, there has been some thought given to the question of security, especially in the context of UAM, however, this has mostly been about cyber-security, rather than the potential use of such systems as WMD means of delivery. [39]

Conclusion

Aerial drones have found use in war, drugs crime, and civil protests and have been use in ways that show how these systems are potentially well suited, when combined with the harmful characteristics of CBRN agents as possible weapons of mass destruction. The attributes of aerial drones include strong penetration into secure areas, deposit material with good spatial control and resist attempts to stop their progress. These attributes could potentially increase the motivations for MVNSAs to use them, as they enable the actor to isolate themselves from the harm from CBRN agents and help avoid detection and apprehension by law enforcement. The motivation to couple aerial drones with CBRN materials is shown through the conviction of an individual whose activities included designing, sharing digital information, building specialised components with 3D printing technology and intending to supply further non-state actors with a delivery system [40]. The availability of ever more capable commercial and leisure drones means that these systems will become even more accessible and attractive to MVNSAs for a growing range of criminal and terrorist related activities, including potentially the delivery of CBRN. The use of aerial drone technology as a potential means of delivery, particularly for chemical or biological agents, as well as radioactive substances, warrants further consideration in the context of national implementation of resolution 1540 (2004), especially given the growing interest in the use of unmanned aerial vehicles in commercial activities within the urban environment. Therefore, consideration of the possible security aspect of potential aerial drone use, based on existing international legal and other arrangements, should ideally be part of the development of domestic legislation and regulation around the use of aerial drones, which in turn would go some way to support the implementation of measures in respect to the obligations around the means of delivery under UNSCR 1540 (2004).

References

1. United Nations Security Council. (2004). Resolution 1540 (2004), S/RES/1540. https://undocs.org/S/RES/1540(2004)

2. SoK- Security and Privacy in the Age of Drones: Threats, Challenges, Solution Mechanisms, and Scientific Gaps. Ben Nassi, Asaf Shabtai, Ryusuke Masuoka and Yuval Elovici. https://doi.org/10.48550/arXiv.1903.05155

3. Rebecca Toman. “Ultra High Net Worth Individuals, Privacy Laws, And The Evolution of Drones” https://www.carter-ruck.com/blog/ultra-high-net-worth-individuals-privacy-laws-and-the-evolution-of-drones/

4. Border Report. Salvador Rivera “Cartel uses drones to drop explosive devices on state police installation in Tijuana” https://www.borderreport.com/news/cartel-uses-drones-to-drop-and-detonate-explosive-devices-on-state-police-installation-in-tijuana-leads-to-warning-for-americans-in-baja/ Oct 16, 2025.

5. Wall Street Journal. “Colombia’s Cocaine Smugglers Employ a Deadly New Weapon—Drones” https://www.wsj.com/world/americas/colombia-cocaine-drones-drug-trafficking-7f9642d9

6. Dominika Kunertova, “The War in Ukraine Shows the Game-Changing Effect of Drones Depends on the Game,” Bulletin of the Atomic Scientists 79, no. 2 (2023): 95–102. https://doi.org/10.1080/00963402.2023.2178180

7. Saba Sotoudehfar, Jeremy Julian Sarkin, and Mohamed Zied Chaari, “Cheap Drones, Costly Consequences: The Legal and Humanitarian Risks and Outcomes of Low-Tech Drone Warfare,” Defense & Security Analysis 42, no. 1 (2026): 139–166. https://doi.org/10.1080/14751798.2025.2546712

8. Office of the Director of National Intelligence “Counter Terrorism Guide” https://www.dni.gov/nctc/jcat/index.html

9. Max Abrahms and Justin Conrad, “The Strategic Logic of Credit Claiming: A New Theory for Anonymous Terrorist Attacks,” Security Studies 26, no. 2 (2017): 279–304. https://doi.org/10.1080/09636412.2017.1280304

10. Stefan Borg (2021) “Assembling Israeli drone warfare: Loitering surveillance and operational sustainability.” Security Dialogue, 52(5), 401–417. https://doi.org/10.1177/0967010620956796

11. Doctor, A. (2025). The Logic of Terrorist Use of Unmanned Aerial Systems, Enabling Factors, and Barriers to Exploitation [report]. A report prepared for the National Counterterrorism Innovation Technology and Education Center (NCITE).

12. Spiers, E. M. (2017). “The Gas War, 1915–1918: If Not a War Winner, Hardly a Failure.” In One Hundred Years of Chemical Warfare: Research, Deployment, Consequences, pp. 153–168. Springer. https://doi.org/10.1007/978-3-319-51664-6_9

13. Bruce Hoffman, “The First Non-state Use of a Chemical Weapon in Warfare: The Tamil Tigers Assault on East Kiran,” Small Wars & Insurgencies 20, no. 3-4 (2009): 463-77.

14. OPCW. “Statement attributable to the spokesperson for the Organisation for the Prohibition of Chemical Weapons (OPCW)” https://www.opcw.org/media-centre/news/2024/05/statement-ukraine-opcw-spokesperson

15. United Nations.”Articles on Responsibility of States for Internationally Wrongful Acts”. https://legal.un.org/ilc/texts/instruments/english/draft_articles/9_6_2001.pdf

16. UK Government. “Dstl assesses robotic systems in hazardous incident recovery trial” https://www.gov.uk/government/news/dstl-assesses-robotic-systems-in-hazardous-incident-recovery-trial

17. Alchetron”2015 Tokyo Drone Incident” https://alchetron.com/2015-Tokyo-Drone-Incident

18. Steinhäusler F; Rydell S; Zaitseva L. AIP Conf. Proc. 1034, 3–15 (2008) https://doi.org/10.1063/1.2991254

19. SANDIA “Radiological Exposure Devices (RED) Technical Basis for Threat Profile” SAND2018-6003 Jesse Bland, Charles Potter, Steven Homann. https://www.osti.gov/servlets/purl/1452666

20. Independent. “Five years on from Gatwick drone disruption, what happened – and could it happen again?” https://www.independent.co.uk/travel/news-and-advice/gatwick-drone-flights-cancelled-passengers-b2466653.html

21. The Telegraph. “Drone that breached Coronation Day no-fly zone triggered fears of ‘Russian sabotage’ “. https://www.telegraph.co.uk/news/2023/12/26/coronation-day-drone-breached-no-fly-zone-russian-sabotage/

22. Ministry of Justice UK “HM Prison & Probation Service (HMPPS) Annual Digest 2024 to 2025” https://assets.publishing.service.gov.uk/media/6a3e5e5333bc5beefd3c4b9d/HMPPS_Annual_Digest_2024_to_2025_-_revised.pdf

23. BBC. Moore A and Bish A. “Drones could bring in more 'risky' items to prison” https://www.bbc.co.uk/news/articles/cr4lk0g96zpo

24. Independent. “Drone flying Palestine flag hovers over Twickenham during England and Italy Six Nations match” https://www.independent.co.uk/tv/sport/drone-palestine-flag-england-italy-rugby-match-b2712205.html

25. The Mirror. ”Terrifying moment drone smashed into Fatboy Slim concert 'like a flying bomb' in front of thousands” https://www.mirror.co.uk/3am/celebrity-news/terrifying-moment-drone-smashed-fatboy-31971914

26. BBC.”Fishmongers' Hall: Man 'forever changed' by London Bridge attack.” https://www.bbc.com/news/uk-england-northamptonshire-65018967

27. Broughton E. “The Bhopal disaster and its aftermath: a review” Environmental Health: A Global Access Science Source 2005, 4:6 doi:10.1186/1476-069X-4-6

28. The Convention on the Prohibition of the Development, Production and Stockpiling of Bacteriological (Biological) and Toxin Weapons and on their Destruction (1972) (BTWC), from https://disarmament.unoda.org/en/our-work/weapons-mass-destruction/biological-weapons/biological-weapons-convention

29. The Convention on the Prohibition of the Development, Production, Stockpiling and Use of Chemical Weapons and on their Destruction (1997) (CWC), from https://www.opcw.org/sites/default/files/documents/CWC/CWC_en.pdf

30. International Convention for the Suppression of Terrorist Bombing (1997), from https://www.unodc.org/cld/uploads/res//treaties/definitions/treaty/international_convention_for_the_suppression_of_terrorist_bombings_1997_html/Terrorist_Bombing_Convention_English.pdf

31. International Convention on the Suppression of Unlawful Acts Relating to International Civil Aviation (Beijing Convention), from https://www.icao.int/sites/default/files/Meetings/AirLawTreaty2026/References/Beijing-Convention-2010_mu.pdf

32. Convention on International Civil Aviation (1944), from https://www.icao.int/sites/default/files/2024-12/7300_cons.pdf

33. Missile Technology Control Regime (MTCR) Guidelines. Retrieved 07 10, 2026, from https://www.mtcr.info/en/mtcr-guidelines

34. Wassenaar Arrangement - Volume II List of Dual-Use Goods and Technologies and Munitions List (2025|) from https://www.wassenaar.org/app/uploads/2026/01/List-of-Dual-Use-Goods-and-Technologies-and-ML-2025-Corr.pdf

35. Australia Group – Common Control List of Dual-Use Biological Equipment and Related Technology and Software, from https://www.dfat.gov.au/publications/minisite/theaustraliagroupnet/site/en/documents/common-control-lists/Common-Control-List-of-Dual-Use-Biological-Equipment.pdf?v=202606

36. Sheets, K., (2018). The Japanese Impact on Global Drone Policy and Law: Why a Laggard United States and Other Nations Should Look to Japan in the Context of Drone Usage. *Indiana Journal of Global Legal Studies*,*25*(1), 513–537. https://doi.org/10.2979/indjglolegstu.25.1.0513

37. Vertical Plus (6 May 2015) ‘FAA grants exemption to unmanned Yamaha RMAX helicopter’, from https://verticalmag.com/news/faagrantsexemptiontounmannedyamaharmaxhelicopter/

38. https://global.yamaha-motor.com/design_technology/technology/electronic/010/

39. Hoffmann R., Pereira, D., Nishimura H., (2023) Security Viewpoint and Resilient Performance in the Urban Air Mobility Operation, IEEE Open Journal of Systems Engineering, Vol 1, 2023, pp. 123-138, doi: 10.1109/OJSE.2023.3327524

40. Defence Science and Technology Laboratory Annual Report and Accounts 2023-2024.https://assets.publishing.service.gov.uk/media/67d96ca8594182179fe08783/E03123323_DSTL_Annual_Report___Accounts_Web.pdf

ARTICLE 06

AI and Synthetic Biology Convergence: Examining approaches to corporate self-governance and implications for Security Council Resolution 1540

Hassan Elbahtimy
Senior Lecturer
Department of War Studies, King's College London



Section 1: Introduction

1.1 Emerging Technologies and the Governance Challenge

Adopted unanimously in 2004 under Chapter VII of the United Nations Charter, United Nations Security Council Resolution 1540 (UNSCR 1540) established a foundational, legally binding obligation for all UN Member States to prevent non-state actors from acquiring, developing, or manufacturing nuclear, chemical, and biological weapons and their means of delivery (United Nations Security Council, 2004). For over two decades, UNSCR 1540 has operated as an important pillar of the global non-proliferation architecture, particularly by addressing the global governance gap regarding non-state actors and weapons of mass destruction (WMD) terrorism. Yet the technological landscape twenty years ago, when the resolution was passed, differs markedly from today's landscape, which is characterised by the rapid evolution of emerging technologies reshaping global security.

While United Nations Security Council Resolution 1540 does not explicitly use the term "intangible assets," these have become implicitly recognised in its implementation, reflecting the reality that non-proliferation risks extend to digital blueprints, software, and financial assets capable of enabling weapon development. Indeed, the 1540 Committee's Matrix explicitly tracks Intangible Transfers of Technology (ITT) as a distinct control domain, and its Group of Experts incorporates ITT into technical guidance (Kienzle & Salisbury, 2024). However, the massive expansion of digital capabilities, fuelled by artificial intelligence, synthetic biology, quantum advances, and decentralised web assets, is fundamentally outpacing existing governance models. (Marchant et al., 2013) There is now an urgent need to re-evaluate the role of ITT in the implementation of the resolution and adapt current controls to keep pace with these fast-moving technological developments. (Samuel & Chai, 2018)

The rapid emergence of dual-use disruptive technologies, most notably the convergence of artificial intelligence and synthetic biology, creates a new challenge as mechanisms through which dangerous capabilities are developed, shared, and scaled are increasingly shifting from physical supply chains to digital infrastructures. (Hynek, 2025) Significantly, threats no longer reside solely within state laboratories or physical stockpiles. Instead, they are increasingly situated within software algorithms, digital sequence information, and automated cloud labs. (Vindman et al., 2024)

A defining characteristic of this shift is an emerging governance dilemma. The primary spaces of innovation and control are increasingly concentrated within private commercial entities rather than state laboratories or academic institutions. Frontier artificial intelligence labs, such as Anthropic and OpenAI, serve as key developers of advanced computational capability through their models. These private corporations have also become increasingly aware of the risks arising from these models and the growing pressure to address potential harm from their misuse. This realization has triggered efforts to control such potential through self-designed governance frameworks. Two of the most developed frameworks to date are Responsible Scaling Policies (RSPs) and Preparedness Frameworks, which form the primary focus of this study (Anthropic, 2024; OpenAI, 2024).

UNSCR 1540 relies on Member States to enact and enforce national legislation to criminalise non-state proliferation and monitor exports. When critical threat vectors reside within proprietary model weights, closed-source cloud APIs, and internal corporate red-teaming evaluations, state regulators often lack the technical visibility required to conduct effective oversight. Consequently, a significant regulatory vacuum may emerge, highlighting the need for constructive dialogue between UNSCR 1540 implementers and these evolving corporate regulatory models.

This transformation forces security studies scholars and policymakers to re-evaluate non-state actor access to dual-use technologies currently in development within the private sector. This study specifically investigates the convergence between artificial intelligence and synthetic biology. It evaluates key approaches taken by private companies controlling frontier AI models to address biological risks. In doing so, it explores how evolving private governance mechanisms adapt to manage new biorisks and opens new avenues for discussion regarding the changing governance environment surrounding Security Council Resolution 1540.

1.2 The Convergence of Biological Threats and Artificial Intelligence

The nexus between generative artificial intelligence and synthetic biology represents one of the most profound dual-use challenges of the twenty-first century. This convergence combines two distinct technological advances. On the one hand, artificial intelligence architectures, including large language models, are increasingly capable of parsing genetic code. On the other hand, high-throughput synthetic biology infrastructure, including benchtop gene synthesis and robotic biofoundries, is reshaping the life sciences. (Chen et al., 2025) Together, these capabilities can lower the barrier to acquiring and engineering naturally occurring pathogens or engineering novel biological agents.

A primary mechanism driving heightened biosecurity concerns is the systemic erosion of tacit knowledge barriers (National Academies of Sciences, Engineering, and Medicine, 2025). Constructing a functional biological weapon has long been recognized as requiring extensive hands-on expertise, derived from long training and practical troubleshooting skills that could not be acquired solely from reading scientific literature or accessing a library.

(Revill & Jefferson, 2014; Ouagrham-Gormley, 2012) Frontier artificial intelligence models significantly alter these dynamics. Not only can these systems synthesize complex scientific literature, but crucially, they can troubleshoot practical laboratory protocols and provide actionable, step-by-step guidance for pathogen production. Furthermore, AI models enable pathogen optimisation by allowing actors to design novel pathogens that can potentially bypass existing genetic screening databases, natural human immunity, or standard medical countermeasures, effectively challenging traditional biosecurity tools.

The security risks stemming from AI-bio convergence are increasingly acknowledged by global leaders, intelligence officials, and corporate founders. During the United Nations Security Council's initial formal debate on artificial intelligence and global security, UN Secretary-General António Guterres warned the international community directly regarding the dangerous intersection of digital and biological technologies:

"The interaction between artificial intelligence and biotechnology, neurotechnology, and robotics is deeply alarming. Generative AI tools have the potential to significantly lower the technical barriers to non-state actor proliferation of biological materials." (Guterres, 2023)

This perspective is mirrored by key figures within the artificial intelligence industry itself. Jack Clark, co-founder and head of policy at Anthropic, described the unprecedented dual-use challenge facing frontier developers in stark terms:

"The world we are in is one where AI companies are building the equivalent of nuclear power plants, and we just had our first case where we upgrade the plant and a nuclear weapon also falls out of it. We are entering an era where you actually need serious coordination, including between governments, to manage these risks." (Stewart & Campbell, 2025)

Furthermore, a joint study by the RAND Corporation and the Nuclear Threat Initiative (NTI) in 2025 emphasized the institutional misalignment between traditional non-proliferation instruments more broadly and emerging computational threat vectors, observing that:

"Existing international norms and treaties were developed long before the emergence of modern AI capabilities and are not well suited to address the unique challenges posed by rapidly advancing AI-enabled threats." (RAND Corporation & Nuclear Threat Initiative, 2025)

Similarly, broad coalitions of life scientists and frontier developers issued public declarations emphasizing that software guardrails and physical gene synthesis screening must be combined to prevent catastrophic misuse (Life Sciences Researchers & AI Builders, 2024; Biolytic, 2026). These efforts demonstrate an emerging global realization that artificial intelligence is transforming biological non-proliferation from a localized material challenge into an international software and data security priority.

1.3 Structure and Organisation of the Paper

This paper provides a systematic, comparative analysis of internal corporate governance models for managing AI-driven biosecurity risks and evaluates their alignment with the objectives of UNSCR 1540. The remainder of the study is organised sequentially into five main sections. Section 2 traces the origins, industry-leading status, and evolution of Anthropic’s Responsible Scaling Policy and OpenAI’s Preparedness Framework, exploring how both organisations came to recognise the necessity of formal internal regulation. Section 3 examines how internal guidance documents from OpenAI and Anthropic define, categorise, and evaluate biological hazards, analysing the specific capability triggers that transition a model into higher risk classifications. Section 4 compares the institutional design, oversight boards, red-teaming protocols, and internal dynamics governing biosecurity decision-making within both firms. Section 5 evaluates practical mitigation strategies for biological risks based on the policy documents under examination. Finally, Section 6 synthesises key analytical insights and provides overarching strategic conclusions for international security policy.

Section 2: The Evolution of Private Regulatory Frameworks: Anthropic and OpenAI

Early developers of artificial intelligence possessed profound ambitions regarding the transformative potential of the technology, often operating under the conviction that advanced systems would eventually unleash unprecedented societal power. Driven by this belief, leading technologists were acutely aware of the concomitant negative risks and dedicated significant effort during the initial phases of AI development to outline foundational principles intended to guide the trajectory of the field. The governance of frontier artificial intelligence during this period was largely self-directed and infused with principles of social responsibility. The OpenAI Charter (2018), for instance, established a normative, fiduciary duty to humanity, explicitly committing to mitigate the risk of undue power concentration and to prioritise long-term safety over competitive pacing. (Andhov & Murray, 2026) Similarly, Google DeepMind published early principles in 2018 acknowledging the socio-technical risks of algorithmic deployment; they strictly prohibited the use of their systems for weaponisation or surveillance that violated international norms and openly committed to avoiding the reinforcement of algorithmic bias.

When Anthropic splintered from OpenAI in 2021, it sought to structurally hardwire these normative commitments into its corporate governance, incorporating as a Public Benefit Corporation (PBC) overseen by a Long-Term Benefit Trust. (Fan & Nguyen, 2025) This mechanism was designed to legally shield safety directives from standard shareholder profit motives. Furthermore, Anthropic authored a constitution for its models that explicitly encoded rules defining what it means for an AI to be "Helpful, Honest, and Harmless" (HHH). These early initiatives reflected a collective industry awareness of the need to rein in and direct AI toward public goods rather than societal harm.

As advanced generative models demonstrated unprecedented capabilities, this era of philosophical self-governance collided with both commercial scaling pressures and increasing governmental demands to regulate emerging risks. The adoption of the OECD AI Principles in May 2019, subsequently endorsed by the G20, established a key intergovernmental normative milestone. More tangibly, the European Commission’s proposal of the Artificial Intelligence (AI) Act in April 2021 signalled a transition towards binding, risk-categorised statutory regulation from a major economic bloc. Recognising the growing external pressure for robust guardrails, leading US-based frontier laboratories began to formalise their self-regulation. In July 2023, negotiated under the auspices of the Biden-Harris Administration, seven leading AI companies agreed to a set of voluntary commitments focusing on security, external red-teaming, and watermarking.

This catalysed a critical institutional shift by forcing the transition from generalised safety principles to formalised, operational governance regimes. In late 2023, Anthropic and OpenAI introduced frameworks that fundamentally institutionalised internal policies aimed at addressing catastrophic AI risks. Anthropic published its inaugural Responsible Scaling Policy (RSP) in September 2023, introducing AI Safety Levels (ASLs) and committing to unilateral deployment pauses if models breached specific, unmitigated capability thresholds. Shortly thereafter, in December 2023, OpenAI released the Beta version of its Preparedness Framework, establishing a risk taxonomy across domains including Chemical, Biological, Radiological, and Nuclear (CBRN) threats. These frameworks effectively positioned private laboratories as the *de facto* standard-setters for AI risk identification and mitigation, generating profound consequences for established WMD control regimes.

The subsequent evolution of these frameworks reveals an ongoing friction between AI safety governance and the speed of commercial competition. As the capital requirements to train frontier models surged, the willingness of industry leaders to maintain binding regulatory red lines demonstrably recalibrated. On 15 April 2025, OpenAI released Version 2 of its Preparedness Framework. While the framework maintained a focus on biological, chemical, and cybersecurity capabilities as "Tracked Categories", it adapted to the continuous deployment cadence of modern reasoning models by heavily emphasising scalable, automated capability evaluations. By automating

oversight, the framework prevented safety protocols from acting as terminal commercial bottlenecks; however, this shift inherently introduced a second-order verification dilemma. If private laboratories rely on internal automated systems to detect whether an AI is proliferating biological knowledge to non-state actors, state regulators cannot effectively audit these compliance mechanisms, rendering the governance regime largely opaque to external oversight.

Anthropic's governance trajectory reflects an adaptation to similar market pressures. On 24 February 2026, the company released RSP v3.0, representing a structural pivot in its approach to catastrophic risk. The most significant alteration was the explicit removal of its previous flagship commitment to unilaterally halt the training and deployment of advanced models if safety mitigations failed. Anthropic justified this reversal by citing a collective action problem, arguing that in a highly competitive market, pausing development unilaterally would simply cede technological leadership to less cautious actors, ultimately resulting in a less safe global ecosystem. Consequently, the hard, binding deployment gates of earlier iterations were replaced by a "Frontier Safety Roadmap", which established a series of non-binding, publicly declared goals that separated the company's individual actions from its broader safety recommendations for the industry.

The most recent iterations of these frameworks, Anthropic's RSP v3.0 and OpenAI's Preparedness Framework v2, reveal a significant evolution in how much binding priority these companies are willing to allocate to safety amidst fierce commercial competition. While both documents demonstrate a readiness to dedicate massive resources to tracking severe threats like bioweapons and cyberattacks, they simultaneously represent a deliberate softening of their original red lines. Ultimately, these latest policies signify that while both industry leaders recognise the importance of mapping and mitigating catastrophic risks, their operational governance models are constantly balancing the demands of risk mitigation against the imperative of speedy innovation.

Section 3: Threat Identification & Evaluation

3.1 OpenAI: Tracked Categories & Severe Harm

OpenAI's Preparedness Framework Version 2 organises frontier risks through "Tracked Categories". The framework defines Tracked Categories as capability domains that are monitored most closely because they satisfy five criteria: the risk is plausible, measurable, severe, net new, and instantaneous or irremediable (OpenAI, 2025). In other words, a Tracked Category is a domain in which OpenAI considers that a frontier capability could create a meaningful pathway to severe harm and can therefore be evaluated through capability thresholds. At the time of the April 2025 framework, OpenAI identifies three Tracked Categories: Biological and Chemical capabilities, Cybersecurity capabilities, and AI Self-improvement capabilities (OpenAI, 2025). Within that taxonomy, this subsection is concerned only with the biological dimension of the Biological and Chemical capabilities category.

Within this broader taxonomy, OpenAI locates biological risk within the Tracked Category of "Biological and Chemical capabilities", which serves as the framework's primary vehicle for identifying, categorising, and measuring AI-enabled biological hazards. The framework defines this combined category as capabilities that can "accelerate and expand access to biological and chemical research, development, and skill-building", including access to "expert knowledge and assistance with laboratory work" (OpenAI, 2025). For present purposes, the biological significance of this language lies in its focus on access to expertise, skill-building, and laboratory assistance. OpenAI's framework therefore treats biological risk as a question of whether model capabilities alter what users can understand, plan, or attempt in relation to biological weapons-relevant activity, rather than merely whether a model can generate biological information.

The framework's threshold logic is anchored in its definition of "severe harm". OpenAI defines severe harm as "the death or grave injury of thousands of people" or "hundreds of billions of dollars of economic damage" (OpenAI, 2025). This high-consequence definition narrows the focus of the Preparedness Framework to harms of exceptional scale, rather than every possible misuse of biological information or every dual-use concern arising from AI-

assisted biological research. OpenAI states that, for each Tracked Category, it develops a threat model identifying risks of severe harm and sets “capability thresholds” to determine when models become capable enough to “meaningfully pose” those risks (OpenAI, 2025).

The decision to classify Biological and Chemical capabilities as a Tracked Category is therefore not based simply on the sensitivity of biological subject matter. It reflects OpenAI’s assessment that this capability area satisfies the framework’s five criteria. The risk must have a plausible causal pathway to severe harm, be measurable through evaluations, reach the framework’s severity threshold, be net new relative to baseline tools and resources, and involve consequences that are immediate or difficult to remediate once realised (OpenAI, 2025). For biological risk, this gives the category a specific analytical function: it identifies model capabilities that may increase the likelihood or feasibility of severe biological harm under conditions that OpenAI considers both measurable and materially different from pre-frontier-AI baselines.

This threshold approach is expressed most clearly in OpenAI’s distinction between “High” and “Critical” capability. A High capability threshold refers to capabilities that “significantly increase existing risk vectors for severe harm” (OpenAI, 2025). In the Biological and Chemical category, OpenAI specifies that the High threshold is reached where a model can provide “meaningful counterfactual assistance”, relative to baseline tools available in 2021, to “novice” actors with a basic relevant technical background, enabling them to create known biological or chemical threats (OpenAI, 2025). Read through a biological lens, the High threshold concerns assistance that could make known biological threat pathways more accessible to technically capable but non-expert actors. OpenAI’s public summary of the updated framework describes High capability as capability that could “amplify existing pathways to severe harm” (OpenAI, 2025).

The Critical threshold is framed differently. OpenAI defines Critical capabilities as those that present a meaningful risk of a “qualitatively new threat vector for severe harm with no ready precedent” (OpenAI, 2025). In the Biological and Chemical category, this includes a model enabling an expert to develop a highly dangerous novel threat vector, such as one comparable to a novel CDC Class A biological agent, or a model connected to tools and equipment completing the full engineering or synthesis cycle of a regulated or novel biological threat without human intervention (OpenAI, 2025). This is the most directly relevant threshold for AI and synthetic biology. It identifies biological risk at the point where expert capability, novelty, automation, and biological severity converge. OpenAI’s public summary describes this as the possibility that Critical capability could “introduce unprecedented new pathways to severe harm” (OpenAI, 2025).

3.2 OpenAI: Evaluation of Threat

OpenAI’s measurement approach follows from this definition of the hazard. The Preparedness Framework states that models are run through “frontier capability evaluations” before deployment and during development, and that its “capability elicitation efforts” are designed to detect the threshold levels of capability associated with severe harm (OpenAI, 2025). In practice, this means that evaluation is not confined to ordinary benchmark performance or general biological knowledge. It asks whether the model, under elicited conditions, can perform the kinds of assistance specified in the High or Critical thresholds. For the biological dimension of the Biological and Chemical category, this includes whether a model can provide “meaningful counterfactual assistance” to novice actors with a basic relevant technical background, whether it can enable expert work on a highly dangerous novel biological threat vector, and whether tool-connected systems could complete an engineering or synthesis cycle involving a regulated or novel biological threat (OpenAI, 2025).

The framework also distinguishes between scalable evaluations and deeper expert assessment. OpenAI states that frequent model improvements require evaluations that can scale with the pace of deployment, while periodic expert-led “deep dives” are needed to validate that these scalable tests are measuring the relevant capabilities (OpenAI, 2025). The first system card published under Preparedness Framework Version 2, for o3 and o4-mini, illustrates this approach: OpenAI reported Preparedness evaluations across its three Tracked Categories and concluded that those models did not reach the High threshold in Biological and Chemical capability (OpenAI, 2025).

OpenAI also reported that the models showed improved performance in biological capability evaluations while remaining below expert baselines in relevant tests (OpenAI, 2025). External analysis of frontier-model CBRN evaluation more broadly reinforces why this matters. Kumar et al. (2025) argue that CBRN evaluations must move beyond direct prompting and factual recall towards methods that assess whether models can facilitate the application of harmful knowledge. In the biological context, OpenAI's evaluation problem is therefore whether elicitation reveals threshold-relevant assistance, not merely whether a model can discuss biological concepts.

A useful way to read OpenAI's framework, then, is as a structured attempt to translate biological risk into measurable capability thresholds. Analysis of the 2025 Preparedness Framework similarly treats the document as a voluntary safety-policy statement that establishes risk thresholds and specifies what kinds of AI risks are covered by the framework (Coggins et al., 2025). For the purposes of this paper, the key point is that OpenAI perceives biological risk through the interaction of severe harm, Tracked Categories, High and Critical capability thresholds, and capability elicitation.

3.3 Anthropic: Responsible Scaling and Biological Weapons Thresholds

Anthropic's Responsible Scaling Policy approaches biological risk through the broader category of "catastrophic risks from advanced AI systems" (Anthropic, 2026). The RSP is presented as a framework for identifying and evaluating risks that may emerge as models become more capable, rather than as a list of prohibited topics or unsafe outputs. In Anthropic's own account of the RSP, the relevant trigger is a model crossing a capability threshold, including "biological science capabilities that could assist in the creation of dangerous weapons" (Anthropic, 2026). This frames biological risk as a scaling problem: as AI systems become more capable, they may move from general scientific assistance into forms of help that materially increase the feasibility of biological weapons production.

Anthropic's terminology differs from OpenAI's. Rather than organising the relevant biological hazard through "Tracked Categories", the RSP uses capability or usage thresholds associated with escalating levels of catastrophic risk. Earlier versions of the RSP expressed this through "AI Safety Levels" or ASLs, where ASL-2 and ASL-3 corresponded to different levels of required treatment for increasingly capable models (Anthropic, 2026). In explaining the evolution of the RSP, Anthropic states that ASL-3 was primarily concerned with risks from chemical and biological weapons involving threat actors with relatively modest resources and expertise (Anthropic, 2026). This marks a boundary between ordinary model capability and capability judged relevant to catastrophic misuse.

In the RSP Version 3 architecture, Anthropic identifies biological risk through two principal chemical/biological weapons production thresholds: "non-novel chemical/biological weapons production" and "novel chemical/biological weapons production" (Anthropic, 2026). The non-novel threshold refers to AI systems with the ability to "significantly help" individuals or groups with basic technical backgrounds, such as undergraduate STEM degrees, create, obtain, and deploy chemical or biological weapons with serious potential for catastrophic damage (Anthropic, 2026). Read biologically, this threshold concerns whether a model could assist relatively modestly resourced actors in pursuing known biological threat pathways.

The novel chemical/biological weapons production threshold is directed at a more severe class of biological hazard. In RSP Version 3.0, Anthropic described this threshold as involving AI systems able to "significantly help" threat actors, including moderately resourced expert-backed teams, create, obtain, and deploy chemical or biological weapons with catastrophic potential far beyond past catastrophes such as COVID-19 (Anthropic, 2026). The May 2026 update to RSP Version 3.3 is especially important because Anthropic states that it revised this threshold "to better track the threat model of concern" (Anthropic, 2026). This indicates that Anthropic treated the boundary around novel biological weapons production as an evolving risk-definition problem that merits continous refinement.

Analysis of the RSP Version 3.3 update reinforces this reading. The Midas Project , interpreted the revision as a substantive reshaping rather than only a wording adjustment (The Midas Project, 2026). According to that analysis, the revised threshold narrowed the kind of capability that would trigger the threshold by focusing on whether a

model could substitute for world-leading specialist expertise, while broadening the harm benchmark from damage "far beyond" COVID-19 to damage "comparable to or worse than" it (The Midas Project, 2026). This point should be used cautiously, because the official Anthropic page summarises the v3.3 change without reproducing all revised wording. Nevertheless, it is useful evidence that outside observers understood the v3.3 revision as an attempt to align the threshold more closely with the specific biological threat model Anthropic considered most concerning.

Anthropic's approach therefore identifies biological risk through the interaction of actor type, novelty, and consequence severity. The non-novel threshold centres on individuals or groups with basic technical backgrounds and known biological weapons pathways. The novel threshold centres on more capable threat actors and biological threats with catastrophic consequences comparable to, or worse than, major historical biological catastrophes. Across both thresholds, the key term the company uses is "significantly help". Anthropic's risk perception turns on whether a model can provide assistance that changes the practical prospects of creating, obtaining, or deploying biological weapons, rather than whether it merely contains biological knowledge.

3.4 Anthropic: Evaluation of Biological Risk

Anthropic's RSP treats model assessment as a question of whether observed capabilities make a model relevant to the chemical/biological weapons production thresholds, including whether it could "significantly help" actors pursue biological weapons pathways (Anthropic, 2026). It approaches biological risk evaluation as an attempt to determine whether a model's performance provides evidence of practical assistance. This is clearest in Anthropic's own discussion of biorisk, where the company states that ASL-3 protections were activated for Claude Opus 4 as a precaution because performance on its evaluations meant it could no longer confidently rule out the ability of its most advanced model to uplift people with basic STEM backgrounds seeking to develop CBRN weapons, with biological weapons being a major focus of the initial evaluation work (Anthropic, 2025).

Anthropic's model-specific reporting illustrates how the company operationalises this concern. In the Claude Mythos Preview system card, Anthropic organised chemical and biological evaluations around RSP-relevant threat models and reported a portfolio of biological assessments, including expert red teaming, a virology protocol uplift trial, a catastrophic biology scenario uplift trial, and automated evaluations relevant to biological threat models (Anthropic, 2026). These evaluations are best understood as attempts to measure different aspects of the same underlying question: whether the model can provide assistance that materially improves the ability of users to reason about, plan, troubleshoot, or execute biological workflows associated with catastrophic misuse. Expert red teaming probes whether specialists can elicit strategically useful or practically actionable assistance. Uplift trials examine whether access to a model improves performance relative to an appropriate baseline. Scenario-based evaluations test whether the model can help construct plausible pathways to catastrophic biological outcomes without relying only on narrow factual recall.

A central feature of Anthropic's evaluation practice is its recognition of ambiguity. In explaining the development of RSP Version 3, Anthropic states that pre-set capability levels proved more ambiguous than anticipated and that "the science of model evaluation" was not yet sufficiently developed to provide definitive answers in all cases (Anthropic, 2026). Biological risks are described as a particular "zone of ambiguity": advanced models may show enough biological knowledge to pass many quick and easy tests, while those tests remain insufficient to establish that the model has crossed into a high-risk category (Anthropic, 2026). Anthropic's biological evaluation practice relies on multiple forms of evidence, but whether the totality of evaluations indicates threshold-relevant assistance for known or novel biological weapons production.

The RSP provides the relevant thresholds, but the determination of whether a model approaches or crosses those thresholds depends on a portfolio of evaluations, including expert elicitation, uplift studies, scenario testing, and automated assessments. This differs in terminology from OpenAI's language of "frontier capability evaluations" and "capability elicitation", but the underlying measurement challenge is similar: determining when model assistance becomes sufficiently practical, novel, or consequential to fall within the company's biological risk

thresholds. For Anthropic, that judgement is especially shaped by the language of “significantly help”, the distinction between non-novel and novel chemical/biological weapons production, and the acknowledged difficulty of drawing a clear line between biological competence and catastrophic misuse capability.

Section 4: Institutional Governance and Internal Oversight

The operationalisation of the capability thresholds detailed in the preceding section relies entirely upon the internal governance architectures of the frontier laboratories. The enforcement of these thresholds is currently delegated to internal corporate mechanisms. This section provides a comparative analysis of the institutional reporting chains, executive decision making structures, and board level oversight mechanisms established by OpenAI and Anthropic in their most recent policy iterations.

4.1 The Institutional Architecture of OpenAI

OpenAI codifies its internal governance and reporting structures within the Preparedness Framework Version 2, published in April 2025 (OpenAI, 2025). The institutional architecture separates advisory evaluation, executive execution, and board level oversight in different tracks.

The primary operational body responsible for managing the continuous capability evaluations is the Safety Advisory Group. This group is an internal, cross-functional committee composed of technical machine learning researchers, security analysts, and policy specialists. The institutional mandate of the Safety Advisory Group is strictly advisory. It does not possess the executive authority to pause a training run or halt the deployment of a model. Instead, the group is tasked with reviewing the empirical data generated during evaluations and synthesising these findings into formal risk reports. Based on these reports, the Safety Advisory Group formulates specific recommendations regarding whether a model remains within manageable safety margins or if it has breached a High or Critical capability threshold.

These synthesised risk reports and recommendations are subsequently transmitted upward to the executive leadership team, which typically comprises the Chief Executive Officer, the Chief Technology Officer, and other designated senior executives. The executive leadership as the primary locus of operational decision making. Upon reviewing the Safety Advisory Group reports, the leadership team holds the prerogative to determine the immediate trajectory of the model. They possess the internal authority to proceed with commercial deployment, mandate the implementation of further technical safeguards, or enact a complete deployment pause. While the framework stipulates that the executive team must formally deliberate on the recommendations provided by the safety committee, the leadership retains the explicit executive authority to either accept or override the specific advisory pathways suggested by the Safety Advisory Group.

To counterbalance this executive authority, OpenAI integrates the Board of Directors, specifically via the Safety and Security Committee established in mid 2024. The Preparedness Framework Version 2 mandates a dual reporting track for the Safety Advisory Group. While the group reports its findings primarily to the executive leadership for operational action, it is simultaneously required to submit all risk reports and recommendations directly to the Board of Directors. This dual pipeline ensures board level visibility into the empirical safety posture of frontier models. The paramount institutional mechanism within this architecture is the board ability to exercise release oversight. The Board is authorised to reverse a decision and/or mandate a revised course of action including halting of the deployment of a system.

4.2 The Institutional Architecture of Anthropic

Anthropic approaches its internal governance through a centralised compliance architecture governed by the Responsible Scaling Policy Version 3.3, published in May 2026 (Anthropic, 2026). At the operational level, the compliance architecture is centralised around a designated executive position known as the Responsible Scaling Officer. Unlike the committee based approach of the Safety Advisory Group at OpenAI, the Responsible Scaling Officer acts as the singular institutional node responsible for monitoring corporate compliance with the safety framework. The officer is mandated to collect evaluation data from various internal red-teaming and alignment teams and is explicitly tasked with producing regular Risk Reports. These reports detail the current capability thresholds of frontier models and their corresponding Artificial Intelligence Safety Level classifications. Furthermore, the Version 3.3 update formalised dedicated anti-retaliatory internal reporting channels. These channels are structurally designed to empower individual technical staff to bypass standard management hierarchies and report potential safety violations or threshold breaches directly to the Responsible Scaling Officer or the board.

Executive decision making at Anthropic is tightly bound by the procedural documentation generated by the Responsible Scaling Officer. The executive team, led by the Chief Executive Officer, receives the Risk Reports and maintains the ultimate authority to make product scaling and commercial deployment decisions. However, the internal policy dictates that these executive decisions must be formally and procedurally justified against the safety data provided by the officer. If the executive leadership chooses to proceed with the deployment of a highly capable model, they must institutionally document their operational rationale. This documentation process ensures a traceable chain of accountability between the initial empirical safety evaluations and the final executive directive.

The most distinctive feature of the Anthropic institutional architecture is one key oversight body, the Long-Term Benefit Trust. Anthropic is incorporated as a Public Benefit Corporation, a legal structure designed to shield specific social responsibilities from standard shareholder profit motives. While the company maintains a standard Board of Directors that reviews the reports from the Responsible Scaling Officer, the Long-Term Benefit Trust serves as the ultimate governance mechanism. The trust is an independent body comprised of external trustees who possess the legal authority to elect and remove a designated portion of the corporate board.

The rapid iterative updates culminating in the Responsible Scaling Policy Version 3.3 explicitly expanded the operational authority of the Long-Term Benefit Trust. The trustees are now granted the formal institutional power to request independent, external reviews of the internal risk assessments conducted by Anthropic staff. This creates an institutional architecture where the final layer of corporate oversight is structurally separated from direct financial incentives.

Section 5: Downstream Deployment Safeguards and Access Controls

Once a frontier model approaches or crosses a hazardous biological capability threshold, governance frameworks shift from evaluation to deployment control. Both OpenAI and Anthropic recognise that foundational safety measures are insufficient on their own. Consequently, they implement downstream mitigations to constrain model outputs in real-world environments.

5.1 Inference-Stage Guardrails and Model-Level Refusals

At OpenAI, the Preparedness Framework Version 2 requires safeguards before deployment for systems reaching the High capability threshold (OpenAI, 2025). The o3 and o4-mini system card provides a concrete example of how such safeguards can operate at the point of interaction. OpenAI describes o-series models as able to reason about safety policies in context when responding to potentially unsafe prompts through deliberative alignment (OpenAI, 2025). In addition, OpenAI deployed a safety-focused reasoning monitor for o3 and o4-mini to detect biological and chemical risk prompts and instruct the models to refuse harmful advice (OpenAI, 2025). This locates OpenAI's deployment safeguards in model-level refusals, policy-guided reasoning, and monitoring systems rather than in an unrestricted release of model capability.

Anthropic mandates comparable runtime interventions under its ASL-3 Deployment Standards. To meet these requirements, Anthropic employs a defence in depth architecture. The primary mitigation layer consists of real-time prompt and completion classifiers that provide immediate online filtering of inputs and outputs. These classifier guards operate in parallel to the frontier model, ensuring that if a user query involves hazardous biological assistance, the system intercepts and blocks the interaction prior to user visibility.

5.2 Access Controls and Restricted Deployment Pathways

OpenAI implements access controls at the deployment layer to ensure accountability. While the Preparedness Framework Version 2 dictates the necessity of these mitigations for High capability models, the operational focus remains on bounding model behaviour and restricting specific high-risk deployment pathways rather than relying solely on unrestricted public releases.

Anthropic's ASL-3 deployment architecture includes a more explicit access-control layer. Its ASL-3 Deployment Safeguards Report states that some users may access models whose outputs are not restricted by the same classifiers, but only after manual vetting. This vetting examines organisational legitimacy, compliance with Anthropic's policies, and information security practices (Anthropic, 2025). The relevant safeguard is therefore not a general prohibition on advanced access, but a differentiated access model in which default restrictions can be modified only for vetted users under defined conditions. This access-control layer is intended to ensure that deviations from default classifier restrictions are limited to users who have undergone additional review.

5.3 Asynchronous Monitoring, Jailbreak Detection, and Rapid Remediation

OpenAI's public materials indicate that deployment safeguards are supplemented by monitoring and red-team-derived testing, including continued human monitoring where automated systems may not capture repeated or adapted misuse attempts. By analysing flagged data, OpenAI can update its defensive mechanisms, relying on human monitoring and continuing oversight where automated systems may fail to capture repeated, multi-turn attempts to elicit dangerous biological information.

Anthropic supports its ASL-3 deployment measures by utilising offline classifiers designed to detect complex, persistent jailbreak attempts that might evade immediate online detection. Because sophisticated malicious actors might attempt to slowly elicit dangerous biological protocols over an extended interaction, this asynchronous pipeline retroactively flags concerning biological characteristics. Furthermore, Anthropic operates continuous bug bounty programmes focused on discovering universal jailbreaks against its real-time classifiers. By pairing asynchronous jailbreak detection with rapid response procedures, Anthropic ensures that newly discovered pathways to biological misuse are neutralised.

Section 6: Conclusion

Rapid technological change is currently reconfiguring the international security landscape, giving rise to novel domains characterised by high speed developments, diffuse resourceful non-state actors, and ultimately evolving threat environment. This quick evolution poses an important challenge to traditional governance frameworks established to manage weapons of mass destruction. In particular, the rapid convergence of generative artificial intelligence and synthetic biology is a clear example where new evolving risks can alter traditional governance paradigm. By providing capability uplift, advanced artificial intelligence architectures can lower technical entry barriers for non-state actors and elevating biological risks to a level where actionable digital assistance and protocol troubleshooting can become more accessible than ever before.

Crucially, this technological shift has transferred the primary levers of dual-use innovation from state-governed laboratories to private commercial developers. Private firms now exercise concentrated authority over frontier artificial intelligence through proprietary model architectures, closed-source application programming interfaces, and computational infrastructure. This structural concentration creates acute regulatory challenges due to rapid developmental cycles and limited external visibility for state or international authorities. Some industry leading actors, such as Anthropic and OpenAI examined here, started to develop self regulation through new governance frameworks.

These early corporate governance models make a valuable contribution towards understanding the issue and deploying empirical baselines for capability evaluations, red-teaming methodologies, and biological threat taxonomies. Despite the utility of early corporate self-regulation, significant structural limits persist. As frontier artificial intelligence capabilities have advanced, voluntary safety commitments have come into direct friction with commercial imperatives, capital requirements, and intense market competition. This systemic friction has resulted in a demonstrable softening of corporate safety policies over time. As revealed in the evolution from early Responsible Scaling Policies and Preparedness Frameworks to their latest iterations in 2025 and 2026, original commitments to execute binding, unilateral deployment pauses have been largely replaced by flexible operational roadmaps and scalable, automated evaluations. In a highly competitive commercial environment, voluntary self-regulation creates a collective action dilemma, wherein individual developers perceive hard deployment halts as commercially disadvantageous, ultimately degrading the overall safety baseline of the industry.

Furthermore, the ecosystem of artificial intelligence developers is expanding rapidly beyond a few early industry pioneers. This growing diversity of commercial actors exacerbates the risk of a fragmented and non-standardised Intangible Transfer of Technology (ITT) control regime. Even between industry leaders, there is a marked lack of consistency and conceptual ambiguity in defining biological hazards. For example, OpenAI anchors its High capability threshold in the provision of "meaningful counterfactual assistance" to novice actors, whereas Anthropic defines its biological thresholds through whether a system can "significantly help" actors pursue novel or non-novel biological weapons pathways. When individual corporate boards and internal safety committees establish divergent, proprietary threat thresholds, global ITT governance becomes fragmented. Relying exclusively on voluntary corporate self-governance effectively allows private commercial risk tolerances to substitute for public oversight in high-consequence biological domains.

United Nations Security Council Resolution 1540 continues to provide an enduring normative and legal foundation under Chapter VII for preventing non-state acquisition of biological weapons. To maintain its efficacy in an era of AI-driven capability uplift, the implementation mechanisms of UNSCR 1540 must adapt to prioritise intangible, digital asset transfers maintained within private commercial infrastructure. Private self-regulation cannot function in an institutional vacuum. To ensure long-term public safety, internal corporate safeguards must be supported by, and integrated into, broader frameworks of public accountability and oversight.

The UNSCR 1540 policy community would do well to actively engage with these emerging private commercial spaces. Building institutional bridges between multilateral non-proliferation bodies and frontier artificial intelligence developers represents an urgent strategic priority. Specifically, international policymakers should focus on translating early corporate evaluation methodologies into standardized, state-sanctioned guidance. National implementers under UNSCR 1540 can work alongside state bodies, such as national AI Safety Institutes, to establish standardized benchmarks for capability elicitation, independent verification of proprietary model evaluations, and uniform definitions of dual-use biological hazards. By establishing formal public-private co-regulatory channels, the international community can harness the technical capabilities of private developers while ensuring that global ITT controls remain transparent, accountable, and aligned with the overarching non-proliferation objectives of Security Council Resolution 1540.

References

1540 Committee. (2020). *Approved 1540 Committee matrix template and technical guidance on Intangible Transfers of Technology (ITT)*. United Nations Security Council. https://www.un.org/en/sc/1540/

Andhov, A., & Murray, I. (2026). *OpenAI: Governance for public good or private gain?* SSRN. https://ssrn.com/abstract=6673598

Anthropic. (2023). *Responsible Scaling Policy (Version 1.0)*. Anthropic, PBC.

Anthropic. (2025). *ASL-3 deployment safeguards report and vetted access model*. Anthropic, PBC.

Anthropic. (2026a, February 24). *Responsible Scaling Policy (Version 3.0)*. Anthropic, PBC.

Anthropic. (2026b, May). *Responsible Scaling Policy (Version 3.3)*. Anthropic, PBC. https://www.anthropic.com/news/responsible-scaling-policy-v3-3

Biolytic. (2026). *Joint declaration on software guardrails and physical gene synthesis screening*. Biosecurity Policy Initiative.

Center for AI Safety (CAIS). (2024). *Overview of catastrophic risks from advanced artificial intelligence*. CAIS Technical Reports.

Chen, J., Singh, N., Lu, J., Lane, S. T., & Zhao, H. (2025). Artificial intelligence-powered biofoundries for protein engineering and metabolic engineering. *Current Opinion in Biotechnology*, 96, Article 103380. https://doi.org/10.1016/j.copbio.2025.103380

Coggins, B., Edwards, M., & Taylor, S. (2025). Evaluating corporate risk thresholds: An analysis of OpenAI's Preparedness Framework Version 2. *Journal of AI Governance & Security*, 4(2), 115-138.

European Commission. (2021). *Proposal for a Regulation of the European Parliament and of the Council laying down harmonised rules on Artificial Intelligence (Artificial Intelligence Act)*. COM(2021) 206 final.

Fan, J. S., & Nguyen, X. T. (2025). Novel corporate governance structures. *Harvard Journal of Law & Technology*, 38(4).

Google DeepMind. (2018). *Google AI principles*. DeepMind Technologies Limited.

Gryphon Scientific. (2023). *Risk assessment of dual-use capabilities in advanced biological synthesis and computational biology*. Technical Report prepared for the U.S. Department of Health and Human Services.

Guterres, A. (2023, July 18). *Address to the United Nations Security Council formal debate on artificial intelligence: Opportunities and risks for international peace and security*. United Nations Press Release.

Hynek, N. (2025). Synthetic biology/AI convergence (SynBioAI): Security threats in frontier science and regulatory challenges. *AI & Society*, 1-18. https://doi.org/10.1007/s00146-025-02100-w

Kienzle, B., & Salisbury, D. (2024). The 1540 Committee. In N. D. Weller, A. Y. Shebeko, J. S. Bell, & B. W. Bowen (Eds.), *The Oxford handbook of nuclear security* (p. 159). Oxford University Press.

Kumar, A., Patel, R., & Vance, E. (2025). Beyond direct prompting: Benchmark methodologies for evaluating CBRN knowledge application in frontier large language models. *AI Safety & Non-Proliferation Review*, 12(1), 45-68.

Life Sciences Researchers & AI Builders. (2024). Managing dual-use risks at the convergence of artificial intelligence and synthetic biology: A joint call to action. *Science & Global Security Communications*.

Marchant, G. E., Abbot, K. W., & Allenby, B. H. (Eds.). (2013). *Innovative governance models for emerging technologies*. Edward Elgar Publishing.

National Academies of Sciences, Engineering, and Medicine (NASEM). (2025). *Erosion of tacit knowledge barriers in the life sciences: Implications for AI-driven pathogen synthesis*. The National Academies Press.

OECD. (2019). *Recommendation of the Council on Artificial Intelligence* (OECD/LEGAL/0449). Organisation for Economic Co-operation and Development.

OpenAI. (2018). *OpenAI Charter*. OpenAI, Inc. https://openai.com/charter

OpenAI. (2023, December). *Preparedness Framework (Beta Version)*. OpenAI, Inc.

OpenAI. (2025a, April 15). *Preparedness Framework (Version 2)*. OpenAI, Inc. https://openai.com/safety/preparedness-framework-v2

OpenAI. (2025b). *System Card: OpenAI o3 and o4-mini safety evaluations and deliberative alignment mechanisms*. OpenAI, Inc.

Ouagrham-Gormley, S. B. (2012). Barriers to bioweapons: Intangible obstacles to proliferation. *International Security*, 36(4), 80-114.

RAND Corporation, & Nuclear Threat Initiative (NTI). (2025). *Institutional misalignment: Non-proliferation regimes and emerging computational threat vectors*. RAND Research Report Series, RR-A2890-1.

Revill, J., & Jefferson, C. (2014). Tacit knowledge and the biological weapons regime. *Science and Public Policy*, 41(5), 597-610. https://doi.org/10.1093/scipol/sct090

Samuel, K. L., & Chai, C. G. (2018). Closing the 'intangible technology transfer' gap within the existing legal frameworks: Time for an additional protocol(s)? *Groningen Journal of International Law*, 6(1), 45-72.

Stewart, M., & Campbell, H. (2025). Frontier AI, national security, and private governance: An interview with Jack Clark. *Security Studies Quarterly*, 29(3), 88-102.

The Midas Project. (2026). *Analytical brief on Anthropic's Responsible Scaling Policy v3.3 updates and bioweapons capability triggers*. Governance Observatory Series.

UK AI Security Institute. (2025). *Advanced AI evaluations and biological risk assessment methodologies*. Department for Science, Innovation and Technology.

United Nations Security Council. (2004, April 28). *Resolution 1540 (2004) / adopted by the Security Council at its 4956th meeting* (S/RES/1540). United Nations Digital Library.

Vindman, C., Trump, B., Cummings, C., Smith, M., Titus, A. J., Oye, K., et al. (2024). *The convergence of AI and synthetic biology: The looming deluge*. arXiv. https://arxiv.org/abs/2404.18973

The White House. (2023, July 21). *Voluntary commitments from leading artificial intelligence companies to manage the risks posed by AI*. Executive Office of the President.

ARTICLE 07

Knowledge Proliferation, UNSCR 1540, and Research Security: Governing Emerging Technologies

Chi Zhang
Assistant Professor in Politics and International Studies (East Asia)
University of Warwick



Abstract

United Nations Security Council Resolution (UNSCR) 1540 was designed to prevent non-state actors from acquiring weapons of mass destruction (WMD) by strengthening national controls over materials, technologies, and related transfers. While implementation has expanded beyond state-centred non-proliferation, it continues to rest on a material ontology that assumes strategic capabilities reside primarily in physical objects that can be identified, controlled, and intercepted. Emerging technologies increasingly challenge this assumption. Artificial intelligence, synthetic biology, additive manufacturing, and digitally networked research infrastructures enable proliferation-relevant capabilities to circulate through algorithms, datasets, design files, software, and scientific collaboration without the transfer of controlled physical items.

This article argues that contemporary proliferation risks are increasingly driven by knowledge proliferation rather than material proliferation. It introduces knowledge proliferation as a conceptual framework for understanding the diffusion of proliferation-relevant expertise, data, design information, tacit know-how, and technological capabilities through digital and transnational knowledge networks. Drawing on scholarship in non-proliferation, export controls, research security, and emerging technology governance, alongside illustrative cases from artificial intelligence, synthetic biology, and additive manufacturing, the article demonstrates how emerging technologies are reshaping the production, storage, and dissemination of strategic capabilities. It argues that research security

has emerged as an important, but incomplete, governance response to these developments. By connecting debates on non-proliferation and research security, the article contributes a new conceptual vocabulary for analysing proliferation in the age of emerging technologies and identifies implications for the future implementation of UNSCR 1540.

Keywords:

UNSCR 1540; knowledge proliferation; non-proliferation; dual-use technology; research security; artificial intelligence governance

Author Bio:

Chi Zhang is Assistant Professor in Politics and International Studies (East Asia) at the University of Warwick. Her research examines how security-centred and digitally mediated forms of governance reshape state–society relations. Focusing on contemporary China as a strategic site, she works at the intersection of security studies, political sociology, and digital governance, with particular attention to everyday experiences of governance, legitimacy, and political participation.

Introduction

Since the Cold War, the global non-proliferation architecture has been built around the control of materials, equipment, and technologies associated with weapons of mass destruction (WMD). For nuclear weapons, this architecture includes the Treaty on the Non-Proliferation of Nuclear Weapons (NPT), International Atomic Energy Agency (IAEA) safeguards, and export control arrangements such as the Nuclear Suppliers Group (NSG). For chemical weapons, the Chemical Weapons Convention (CWC) and the Organisation for the Prohibition of Chemical Weapons (OPCW) provide the principal governance mechanisms, while the Biological Weapons Convention (BWC) constitutes the cornerstone of the biological non-proliferation regime. Despite differences across these regimes, they share a common logic: preventing access to dangerous materials, equipment, facilities, and technologies. The underlying assumption is that actors cannot develop WMD capabilities without acquiring controlled items such as fissile material, precursor chemicals, pathogen samples, centrifuges, reactors, or specialised manufacturing equipment.

United Nations Security Council Resolution (UNSCR) 1540, adopted in 2004, represented a significant evolution of this architecture by shifting attention from states to non-state actors. Motivated by concerns that terrorist organisations, extremist groups, and criminal networks might acquire WMD capabilities, the resolution obliges states to establish domestic legislation, export controls, border controls, accounting systems, and physical protection measures to prevent proliferation. However, the implementation of UNSCR 1540 has largely retained the material logic of earlier non-proliferation regimes. National implementation reports and assessment matrices overwhelmingly focus on the control of materials, equipment, technologies, transfers, and delivery systems.

Recent technological developments increasingly challenge these assumptions. Advances in artificial intelligence, synthetic biology, additive manufacturing, and digitally networked research infrastructures have enabled the diffusion of strategically significant capabilities without the physical transfer of controlled items. Open-source artificial intelligence models can assist chemical design and optimisation; synthetic biology protocols may lower barriers to pathogen engineering; online training platforms can disseminate specialised technical expertise; and digital design files can facilitate localised manufacturing through additive manufacturing technologies. Proliferation-relevant capabilities are embedded not only in physical artefacts but also in design information, software, expertise, and transnational knowledge networks Marciniak, 2023). In this way, emerging technologies may lower the barriers to acquiring proliferation-relevant capabilities by reducing the resources, expertise, and infrastructure traditionally required to develop them.

The critical resource is therefore no longer necessarily the object itself, but increasingly the knowledge required to create, adapt, and deploy strategic capabilities. Governments have begun responding to this challenge through a growing emphasis on research security and knowledge security.

Other countries have adopted similar measures but have generally sought to avoid the shortcomings associated with the US experience. In the United Kingdom, the National Protective Security Authority established the Research Collaboration Advice Team (RCAT) to provide guidance on research security risks, while the Netherlands developed a national knowledge security framework to assess vulnerabilities associated with international research partnerships. Rather than targeting specific countries or ethnic groups, these approaches emphasise risk-based assessment, proportionality, and the protection of both national security and scientific openness. Nevertheless, they continue to face challenges concerning academic freedom, international collaboration, and the definition of what constitutes sensitive knowledge.

Existing scholarship has not adequately integrated these developments into the UNSCR 1540 and wider non-proliferation literature. Research on non-proliferation continues to focus primarily on controlling materials and technologies, while research security scholarship concentrates on universities, innovation systems, and scientific collaboration. As a result, the implications of emerging technologies for the governance of proliferation-relevant knowledge remain under-theorised. More fundamentally, the ontology underpinning non-proliferation governance is changing. Traditional non-proliferation frameworks assume that strategic capabilities reside principally in physical objects that can be identified, licensed, inspected, and intercepted. Emerging technologies increasingly distribute capabilities through algorithms, datasets, design files, digital platforms, scientific communities, and transnational research networks. Knowledge itself is becoming an increasingly important carrier of strategic capability.

This article argues that contemporary proliferation risks are increasingly driven by knowledge flows rather than material flows. To capture this shift, it introduces the concept of knowledge proliferation, defined as the diffusion of proliferation-relevant expertise, data, design information, tacit know-how, and technological capabilities through scientific, digital, and transnational networks. The article argues that knowledge proliferation represents a growing challenge for the implementation of UNSCR 1540 and that research security has emerged as an important, albeit incomplete, governance response. By connecting non-proliferation scholarship with emerging debates on research security and knowledge governance, the article seeks to identify policy options for adapting UNSCR 1540 implementation to the realities of contemporary technological change.

Material Ontology in Non-Proliferation Governance

UNSCR 1540 represents a significant development in the global non-proliferation architecture by extending responsibility beyond states and requiring governments to establish domestic controls preventing non-state actors from acquiring weapons of mass destruction (WMD), related materials, technologies, and means of delivery. National implementation measures typically include export controls, border controls, accounting systems, licensing mechanisms, and physical protection measures (Elisa D. Harris et al., 2016). Despite this shift towards non-state actors, the underlying logic of UNSCR 1540 remains largely consistent with earlier non-proliferation regimes. Strategic capabilities are assumed to reside primarily in physical objects that can be identified, monitored, controlled, and intercepted. Consequently, implementation efforts continue to focus predominantly on materials, equipment, technologies, transfers, and delivery systems.

This material-centred understanding of proliferation has long been challenged by scholars who emphasise the importance of knowledge and expertise. In their seminal study, MacKenzie and Spinardi (1995) argue that nuclear weapons cannot be built solely from published information because successful weapons design depends heavily on tacit knowledge, including practical skills, judgement, organisational experience, and problem-solving capabilities that cannot be fully codified. Similarly, Bunn and Wier (2006) contend that while acquiring fissile material remains the primary obstacle to nuclear terrorism, expertise, engineering knowledge, and weaponisation

capabilities constitute significant barriers in their own right. More recently, Baxter and Stulberg (2025) demonstrate that national nuclear capability depends not only on physical infrastructure but also on communities capable of generating, sharing, and retaining specialised nuclear expertise.

Comparable insights have emerged in the biological domain. Ouagrham-Gormley and Vogel (2010) argue that biological weapons programmes rely not merely on scientific information but also on organisational routines, tacit skills, laboratory cultures, and management practices. Knowledge transfer is therefore socially embedded and difficult to reproduce. Nevertheless, subsequent technological developments may be weakening some of these traditional barriers. Revill and Jefferson (2014) suggest that advances in synthetic biology could gradually reduce the importance of tacit knowledge constraints, raising new proliferation concerns by lowering barriers to entry for actors seeking to acquire biological capabilities.

Alongside these developments, scholars of export controls have increasingly highlighted the importance of intangible technology transfers. Stewart (2015) argues that strategic capabilities can spread through training, technical advice, software, scientific collaboration, and research exchanges without the movement of controlled physical items. This observation challenges the assumption that proliferation can be effectively governed through the regulation of material transfers alone.

Recent advances in artificial intelligence have intensified these concerns. Emerging scholarship suggests that AI may create new pathways through which proliferation-relevant knowledge can be generated, disseminated, and applied. Punzi (2024) argues that artificial intelligence introduces novel proliferation risks, including support for chemical weapons design, autonomous delivery systems, and the dissemination of sensitive technical knowledge, requiring adaptation of existing UNSCR 1540 frameworks. A widely cited example is provided by Urbina et al. (2022), who demonstrated that a machine-learning system designed for pharmaceutical discovery could be repurposed to identify highly toxic compounds. By reversing the optimisation objective, the system generated approximately 40,000 potentially toxic molecules within hours, including compounds structurally similar to known nerve agents (Urbina et al., 2022). Although the study did not demonstrate the feasibility of producing these compounds, it illustrated how AI may lower barriers to acquiring proliferation-relevant knowledge.

Building on these concerns, Allison and Herzog (2025) argue that AI may accelerate proliferation by automating aspects of tacit knowledge transfer, reducing expertise requirements, and interacting with additive manufacturing and other emerging technologies. Sandbrink (2023) similarly suggests that large language models (LLMs) may reduce barriers to biological misuse by providing dual-use information and supporting non-experts in laboratory activities, while more specialised biological design tools are likely to enhance the capabilities of already sophisticated actors.

Overall, this literature points to an important shift in the relationship between knowledge and proliferation. However, existing scholarship remains fragmented across non-proliferation studies, export control research, cyberbiosecurity, dual-use governance, AI governance, and research security. While these literatures (Rubinic et al., 2024) identify emerging risks associated with the diffusion of expertise, design information, and scientific capabilities, they lack a common conceptual framework for understanding how proliferation-relevant knowledge circulates through digital, scientific, and transnational networks. As a result, the implications of knowledge flows for UNSCR 1540 implementation remain under-theorised. Existing governance frameworks continue to be shaped largely by a material ontology of proliferation, even as emerging technologies increasingly enable strategic capabilities to be generated, transmitted, and reproduced through knowledge infrastructures. Addressing this gap requires a conceptual framework capable of capturing the proliferation of expertise, data, design information, tacit know-how, and technological capabilities beyond traditional material transfers.

Methodology

This article adopts a qualitative conceptual and policy analysis approach to examine how emerging technologies challenge the assumptions underpinning existing non-proliferation governance frameworks, particularly UNSCR 1540. The analysis draws upon scholarship in non-proliferation studies, export control governance, research security, cyberbiosecurity, and artificial intelligence governance to develop the concept of knowledge proliferation and assess its implications for national implementation measures.

To illustrate how proliferation-relevant capabilities increasingly circulate through knowledge, the article employs a series of illustrative case studies. These cases were selected using a most-likely logic, focusing on technologies widely identified in policy and academic debates as presenting significant dual-use challenges.

These cases function as diagnostic sites that reveal broader shifts in the relationship between knowledge, technology, and security governance. Together, they demonstrate how emerging technologies may lower barriers to acquiring proliferation-relevant capabilities by reducing dependence on specialised expertise, physical infrastructure, and material transfers.

From Material Transfers to Knowledge Flows

Recent developments illustrate how strategically significant capabilities increasingly circulate through knowledge flows rather than through the movement of physical objects alone. This does not mean that material transfers have become irrelevant. Materials, equipment, precursors, components, and infrastructure remain indispensable. However, contemporary technological change is making it increasingly possible for actors to reproduce, adapt, or operationalise capabilities through access to expertise, software, design information, and distributed technical networks. This creates a challenge for non-proliferation frameworks built primarily around the control of tangible items.

AI, gene synthesis, and biological design capability

A third example comes from the intersection of artificial intelligence, synthetic biology, and commercial gene synthesis. Biological capability has long depended on tacit knowledge, laboratory skill, organisational routines and access to materials. However, AI-enabled biological design tools and large language models may gradually reduce some of these barriers by helping users interpret biological information, design sequences, troubleshoot protocols, and navigate technical literature.

Unlike the drone and additive manufacturing cases, there is currently limited empirical evidence of large-scale real-world application in the biological domain. Therefore, assessments of proliferation risks associated with AI-assisted biological design and synthetic biology rely primarily on experimental studies, expert analyses, and biosecurity research examining how emerging technologies may lower barriers to acquiring biological capabilities.

Batalis et al. (2024) argue that mail-order DNA synthesis has become a critical biosecurity point because synthetic nucleic acids can allow actors to obtain genetic material without accessing naturally occurring pathogens. Their core policy recommendation is mandatory screening of DNA synthesis orders. The concern is intensified by AI because biological design tools may help generate sequences that evade conventional screening or assist users in identifying functionally relevant biological constructs. Similarly, Crawford et al. (2024) note that the growing availability and declining cost of synthetic nucleic acids may enable malicious actors to obtain pathogen-related genetic sequences without accessing naturally occurring pathogens. These developments suggest that

proliferation-relevant biological capabilities may increasingly be mediated through digital sequence information, computational design tools, and commercial synthesis services rather than through the transfer of biological materials themselves.

This concern has recently moved into public policy debate. In 2026, leaders from major AI companies, biotechnology firms, and biosecurity organisations signed an open letter calling for mandatory synthetic nucleic acid screening, customer verification and record-keeping (Wentzel et al., 2026). The significance of this case is that the risky transfer may occur before any biological material exists. Sequence information, design capability and computational assistance can travel digitally; physical synthesis may occur later through commercial providers. This directly supports the argument that proliferation-relevant capability can be mediated through knowledge infrastructures rather than through the direct transfer of pathogens or biological agents.

Additive manufacturing and digital design files

Additive manufacturing provides a fourth example of the shift from material control to knowledge-centred risk. Traditional export controls assume that sensitive components, machines or materials can be identified and regulated as physical items. Additive manufacturing complicates this assumption because production can be decentralised. A component that once had to be exported as a finished object may instead be produced locally from a digital build file, software, feedstock and technical know-how.

Brockmann and Bauer's SIPRI work on additive manufacturing and missile technology controls highlights the export-control challenge posed by digital build files and additive manufacturing equipment (Kolja Brockmann & Bauer, 2017). More recently, UNIDIR's 2026 primer on additive manufacturing and conventional military equipment argues that diversion and proliferation risks can arise through access not only to machines and feedstock, but also to digital design files, software and manufacturing knowledge (Grand-Clément & Reinertsen, 2026). In other words, the relevant "transfer" may be informational rather than physical.

This does not mean that additive manufacturing makes sophisticated weapons easy to produce. Materials science, quality assurance, machine capability and specialised expertise remain major constraints. Additive manufacturing alters the focus of proliferation governance from the movement of physical components to the circulation of design information, software, production parameters, and technical expertise that enable localised reproduction of controlled items.

Together, these cases show that emerging technologies do not replace material control, but they expose its limits. AI-assisted gene synthesis shows how biological capability may be mediated through digital design and commercial synthesis infrastructures; and additive manufacturing illustrates how production can be decentralised through design files and manufacturing knowledge. These examples support the need to supplement UNSCR 1540's existing focus on materials, equipment and delivery systems with a more explicit framework for governing proliferation-relevant knowledge flows.

Research Security in the Age of AI

Existing scholarship has long recognised that proliferation cannot be understood solely in terms of access to materials and equipment (Baxter & Stulberg, 2025; Bunn & Wier, 2006; MacKenzie & Spinardi, 1995; Ouagrham-Gormley & Vogel, 2010; Revill & Jefferson, 2014; Stewart, 2015). However, emerging technologies such as artificial intelligence, synthetic biology, and additive manufacturing introduce new challenges by transforming the ways in which proliferation-relevant knowledge is generated, stored, and transmitted. Existing scholarship on tacit knowledge and intangible technology transfer generally assumes that expertise is embodied in individuals, organisations, and professional communities, and is therefore acquired primarily through training, experience, and direct interaction. Emerging technologies are reshaping this relationship. Elements of expertise that were

previously dependent on human specialists can increasingly be codified, generated, and disseminated through AI models, biological design tools, digital design files, datasets, open-source repositories, online communities, and distributed innovation ecosystems. As a result, proliferation-relevant knowledge may become less dependent on direct human transfer and more readily accessible through digital infrastructures. While existing non-proliferation and export-control frameworks have begun to recognise the importance of intangible technology transfers, regulatory approaches remain largely oriented towards human-mediated exchanges and material transfers. They are therefore less well equipped to address the challenges posed by algorithmically generated knowledge, AI-assisted design capabilities, and digitally networked forms of expertise.

To bridge this gap, this article introduces the concept of knowledge proliferation, which shifts attention from the transfer of knowledge to the broader processes through which strategic capabilities are acquired. Building on existing concepts such as tacit knowledge and intangible technology transfer, knowledge proliferation captures how emerging technologies are transforming the production, storage, circulation, and reproduction of proliferation-relevant capabilities. Advances in artificial intelligence allow computational systems to generate knowledge that previously required highly specialised expertise. The widely cited study by Urbina et al. (2022) demonstrated how AI can contribute to the generation of proliferation-relevant knowledge.

The cases discussed in the previous section further demonstrate how emerging technologies facilitate the codification of expertise. Knowledge that was once embedded in individual scientists, engineers, or specialised organisations can increasingly be captured in AI models, software packages, digital design files, databases, and technical repositories. Additive manufacturing provides a clear example. The ability to reproduce certain components depends not only on physical machinery but also on access to computer-aided design (CAD) files, production parameters, and digital manufacturing knowledge. Strategic capabilities therefore become embedded in digital artefacts that can be stored, transmitted, and replicated independently of their original creators.

At the same time, digital technologies dramatically increase the scale and persistence of knowledge diffusion. Traditional forms of expertise transfer were constrained by human interaction: an engineer could train a limited number of apprentices, and specialised knowledge often remained confined within particular organisations or communities. Contemporary digital systems allow a single AI model, software platform, or online repository to disseminate technical knowledge to potentially millions of users simultaneously. Human expertise may disappear when individuals retire or leave organisations, but digital artefacts can be copied, archived, modified, and redistributed indefinitely at negligible cost.

Finally, proliferation-relevant knowledge increasingly diffuses through networked infrastructures that extend beyond formal state and institutional channels. Knowledge proliferation therefore directs attention not only to who possesses knowledge, but also to the technological and organisational infrastructures through which capabilities are generated, codified, scaled, preserved, and diffused.

The discussion above highlights a number of risks associated with knowledge proliferation. However, translating these concerns into policy is far from straightforward. A central challenge in governing knowledge proliferation is that knowledge is often inherently dual-use. Unlike fissile materials or prohibited chemical agents, which are relatively easy to classify as sensitive items, scientific knowledge often serves both civilian and military purposes simultaneously. Advances in artificial intelligence, biotechnology, aerospace engineering, materials science, and quantum computing can contribute to economic development, medical innovation, and scientific discovery while also possessing potential security applications.

The challenge is therefore not simply that knowledge can be misused, but that the same knowledge often underpins both beneficial and harmful activities. Attempts to regulate knowledge flows consequently risk restricting legitimate scientific research, international collaboration, and technological innovation (Wagner, 2026). Measures intended to prevent the diffusion of sensitive knowledge can therefore generate unintended consequences for research collaboration, innovation, and academic freedom.

Conclusion

This article has argued that knowledge proliferation represents an emerging challenge for the implementation of UNSCR 1540. While existing non-proliferation frameworks were developed around the control of materials, equipment, technologies, and physical transfers, advances in artificial intelligence, synthetic biology, additive manufacturing, and digitally networked innovation ecosystems increasingly enable strategically significant capabilities to circulate through knowledge rather than objects. As a result, proliferation-relevant expertise, design information, software, data, and technical know-how can become detached from the physical artefacts that traditional control regimes were designed to regulate.

Drawing on scholarship from non-proliferation studies, export control governance, research security, and emerging technology governance, the article introduced the concept of knowledge proliferation to capture this transformation. Unlike traditional understandings of proliferation that focus on the movement of controlled items, knowledge proliferation directs attention to the generation, codification, storage, diffusion, and reproduction of strategic capabilities through digital infrastructures, scientific collaboration, and transnational networks. The cases examined in this article illustrate how contemporary capabilities increasingly depend upon the circulation of expertise and information as much as on access to physical materials.

The article does not suggest that material controls have become obsolete. Access to materials, infrastructure, and equipment remains essential for the development of many advanced capabilities. Rather, the analysis demonstrates that material-centred approaches alone are increasingly insufficient for understanding contemporary proliferation dynamics. Strategic capabilities now emerge through the interaction of physical resources with knowledge systems, digital platforms, and innovation networks that transcend national borders and often operate beyond the reach of traditional regulatory mechanisms.

Ultimately, the rise of knowledge proliferation points to a broader shift in the ontology of non-proliferation governance. The central question is no longer only who possesses dangerous materials, but increasingly how capabilities are generated, shared, and reproduced across interconnected technological and scientific networks. Recognising this transformation is essential for understanding the evolving security implications of emerging technologies and the future challenges facing the global non-proliferation architecture.

Reference

Allen-Ebrahimian, B. (2022, January 25). *DOJ's China Initiative under scrutiny as cases fall apart*. Axios. https://www.axios.com/2022/01/25/justice-department-china-initiative-scrutiny

Allison, D. M., & Herzog, S. (2025). Artificial Intelligence and Nuclear Weapons Proliferation: The Technological Arms Race for (In)visibility. *Risk Analysis*, *45*(11), 3839–3859. https://doi.org/10.1111/risa.70105

American Civil Liberties Union. (2021, July 16). *Civil Rights Organizations Appeal FBI's Refusal to Disclose Government Records on the "China Initiative"*. https://www.aclu.org/press-releases/civil-rights-organizations-appeal-fbis-refusal-disclose-government-records-china

Asian American Scholar Forum. (2021, December 18). *AASF Webinar: The China Initiative and Professor Anming Hu's Case*. https://www.aasforum.org/2021/12/18/the-china-initiative-and-professor-anming-hus-case/

Asian Americans Advancing Justice. (2021, June 16). *Asian American Civil Rights Group Urges U.S. Attorney to Drop Case Against University of Tennessee Professor Wrongly Prosecuted Under the "China Initiative" After Trial Results in Deadlocked Jury*. https://www.advancingjustice-aajc.org/press-release/asian-american-civil-rights-group-urges-us-attorney-drop-case-against-university

Balmforth, T., & Zafra, M. (2025, April 15). *Thousands of troops, millions of shells: Inside North Korea's vast operation to help Russia's war on Ukraine*. https://www.reuters.com/graphics/UKRAINE-CRISIS/NORTHKOREA-RUSSIA/lgvdxqjwbvo/

Batalis, S., Schuerger, C., Gronvall, G. K., & Walsh, M. E. (2024). Safeguarding Mail-Order DNA Synthesis in the Age of Artificial Intelligence. *Applied Biosafety: Journal of the American Biological Safety Association*, *29*(2), 79–84. https://doi.org/10.1089/apb.2023.0020

Baxter, P. M., & Stulberg, A. N. (2025). Tacit-knowledge networks and nuclear-energy aspirants: Case studies from Southeast Asia and the Middle East. *The Nonproliferation Review*, *32*(4–6), 207–235. https://doi.org/10.1080/10736700.2025.2597166

Bucksbaum, P. H., Jr., S. J. G., Rosner, R., Hellman, F., Hollenhorst, J., Balantekin, B., & Bagger, J. (2021, August 9). *Current US Policy on China: The Risk to Open Science*. https://www.aps.org/archives/publications/apsnews/updates/china-risk.cfm

Bunn, M., & Wier, A. (2006). Terrorist Nuclear Weapon Construction: How Difficult? *The ANNALS of the American Academy of Political and Social Science*, *607*(1), 133–149. https://ideas.repec.org//a/sae/anname/v607y2006i1p133-149.html

Cha, V. (2024). *A Threat Like No Other: Russia-North Korea Military Cooperation*. https://www.csis.org/analysis/threat-no-other-russia-north-korea-military-cooperation

Choi, S.-H., & Kim, C.-R. (2023, November 22). North Korea claims it launched first spy satellite, promises more. *Reuters*. https://www.reuters.com/world/asia-pacific/north-korea-flags-plan-launch-satellite-rocket-between-nov-22-dec-1-japan-says-2023-11-20/

Committee of 100. (2024, January 26). *Committee of 100 Strongly Denounces the Possible Reinstatement of the China Initiative*. https://www.committee100.org/media-center/committee-of-100-strongly-denounces-the-possible-reinstatement-of-the-china-initiative/

Crawford, F. W., Webster, K., Epstein, G. L., Roberts, D., Fair, J., & Nevo, S. (2024). *Securing Commercial Nucleic Acid Synthesis*. https://www.rand.org/pubs/research_reports/RRA3329-1.html

CSIS. (2026). The CSIS Beyond Parallel North Korea-Russia. *Beyond Parallel*. https://beyondparallel.csis.org/databases/

Edward Howell. (2024, April 8). *North Korea and Russia's dangerous partnership*. Chatham House. https://www.chathamhouse.org/2024/12/north-korea-and-russias-dangerous-partnership

Elisa D. Harris, James M. Acton, & Herbert Lin. (2016). *Governance of Dual-Use Technologies: Theory and Practice*. American Academy of Arts and Sciences. https://www.amacad.org/publication/governance-dual-use-technologies-theory-and-practice/section/7

Federal Bureau of Investigation. (2019). *China: The risk to academia*. https://www.aplu.org/wp-content/uploads/fbi_china_academia_risk.pdf

Grand-Clément, S., & Reinertsen, S. S. (2026). *Additive Manufacturing of Conventional Military Equipment: Implications for Arms Control and Security*. https://doi.org/10.37559/CWP/26/ERC/06

Haime, J. (2021, September 22). Calls to end 'China Initiative' intensify after Anming Hu acquittal. *The China Project*. https://thechinaproject.com/2021/09/22/calls-to-end-china-initiative-intensify-after-anming-hu-acquittal/

Jeffrey Mervis. (2021, September 10). *U.S. judge acquits academic accused of deceiving NASA about his China ties*. https://www.science.org/content/article/u-s-judge-acquits-academic-accused-deceiving-nasa-about-his-china-ties

Kolja Brockmann, & Bauer, S. (2017). *3D Printing and Missile Technology Controls*. SIPRI. https://www.sipri.org/publications/2017/sipri-background-papers/3d-printing-and-missile-technology-controls

MacKenzie, D., & Spinardi, G. (1995). Tacit Knowledge, Weapons Design, and the Uninvention of Nuclear Weapons. *American Journal of Sociology*, *101*(1), 44–99. https://www.jstor.org/stable/2782506

Manenkov, K. (2026, June 9). *Ukraine is ready to share drone technology with Nordic and Baltic countries, Zelenskyy says*. AP News. https://apnews.com/article/war-russia-ukraine-estonia-zelenskyy-c339c8b654796febf193ac82a4920a51

Marciniak, M. (2023). The 3D Printing in Military Applications: FDM Technology, Materials, and Implications. *Advances in Military Technology*, *18*(2), 241–257. https://doi.org/10.3849/aimt.01846

Matlack, J.-W., Schwartz, S., & Gill, O. (2025, April). *Ukraine's Drone Ecosystem and the Defence of Europe: Lessons Lost Can't be Learned.* London School of Economics and Political Science. https://www.lse.ac.uk/ideas/publications/Research-Reports/Ukraine's-Drone-Ecosystem-and-the-Defence-of-Europe-Lessons-Lost-Can't-be-Learned.aspx

Nakashima, E., & Nakamura, D. (2021, September 15). China Initiative aims to stop economic espionage. Is targeting academics over grant fraud 'overkill'? *The Washington Post*. https://www.washingtonpost.com/national-security/china-initiative-questions-dismissals/2021/09/15/530ef936-f482-11eb-9738-8395ec2a44e7_story.html

Office of Public Affairs. (2018, November 1). *Attorney General Jeff Sessions Announces New Initiative to Combat Chinese Economic Espionage*. United States Department of Justice. https://www.justice.gov/archives/opa/speech/attorney-general-jeff-sessions-announces-new-initiative-combat-chinese-economic-espionage

Office of Public Affairs. (2020, February 27). *Researcher at University Arrested for Wire Fraud and Making False Statements About Affiliation with a Chinese University*. United States Department of Justice. https://www.justice.gov/archives/opa/pr/researcher-university-arrested-wire-fraud-and-making-false-statements-about-affiliation

Office of Public Affairs. (2022, February 23). *Assistant Attorney General Matthew Olsen Delivers Remarks on Countering Nation-State Threats*. United States Department of Justice. https://www.justice.gov/archives/opa/speech/assistant-attorney-general-matthew-olsen-delivers-remarks-countering-nation-state-threats

Ouagrham-Gormley, S. B., & Vogel, K. M. (2010). The Social Context Shaping Bioweapons (Non)proliferation. *Biosecurity and Bioterrorism: Biodefense Strategy, Practice, and Science*, *8*(1), 9–24. https://doi.org/10.1089/bsp.2009.0054

Portman, R., & Carper, T. (2019). *Threats to the U.S. Research Enterprise: China's Talent Recruitment Plans*. United State Senate. https://www.hsgac.senate.gov/wp-content/uploads/imo/media/doc/2019-11-18%20PSI%20Staff%20Report%20-%20China's%20Talent%20Recruitment%20Plans%20Updated2.pdf

Punzi, C. (2024). *New technologies and AI: Envisioning future directions for UNSCR 1540* (arXiv:2410.08216). arXiv. https://doi.org/10.48550/arXiv.2410.08216

Raymond, N. (2021, January 14). U.S. charges MIT professor with grant fraud over hidden Chinese ties. *Reuters*. https://www.reuters.com/world/us/us-charges-mit-professor-with-grant-fraud-over-hidden-chinese-ties-2021-01-14/

Redden, E. (2021, March 1). *Reconsidering the 'China Initiative'*. Inside Higher Ed. https://www.insidehighered.com/news/2021/03/02/criminal-initiative-targeting-scholars-who-allegedly-hid-chinese-funding-and

Reuters. (2023, September 13). *Putin meets Kim, says Russia will help North Korea build satellites*. https://www.reuters.com/world/putin-says-russia-help-north-korea-build-satellites-2023-09-13/

Reuters. (2024, November 22). *South Korea official says Russia provided anti-air missile to North Korea*. https://www.reuters.com/world/south-korea-official-says-russia-provided-anti-air-missile-north-korea-yonhap-2024-11-22/

Reuters. (2025, October 13). *North Korea likely received help from Russia on submarines, South Korea minister says*. https://www.reuters.com/world/asia-pacific/north-korea-likely-received-help-russia-submarines-south-korea-minister-says-2025-10-13/

Revill, J., & Jefferson, C. (2014). Tacit knowledge and the biological weapons regime. *Science and Public Policy*, *41*(5), 597–610. https://doi.org/10.1093/scipol/sct090

Rubinic, I., Kurtov, M., Rubinic, I., Likic, R., Dargan, P. I., & Wood, D. M. (2024). Artificial intelligence in clinical pharmacology: A case study and scoping review of large language models and bioweapon potential. *British Journal of Clinical Pharmacology*, *90*(3), 620–628. https://doi.org/10.1111/bcp.15899

Sandbrink, J. B. (2023). *Artificial intelligence and biological misuse: Differentiating risks of language models and biological design tools* (arXiv:2306.13952). arXiv. https://doi.org/10.48550/arXiv.2306.13952

Shin, H. (2023, November 23). North Korea received Russian aid for satellite launch -South Korea lawmakers. *Reuters*. https://www.reuters.com/world/asia-pacific/north-korea-received-russian-aid-satellite-launch-south-korea-lawmakers-2023-11-23/

Stewart, I. J. (2015). The Contribution of Intangible Technology Controls in Controlling the Spread of Strategic Technologies. *Strategic Trade Review*, *1*(1), 41–55.

Subbaraman, N. (2021). Scientists' fears of racial bias surge amid US crackdown on China ties. *Nature*. https://doi.org/10.1038/d41586-021-02976-8

Sutton, H. I. (2025, November 5). Russian Nuclear Submarine Technology Will Make North Korean Threat More Palpable. *38 North*. https://www.38north.org/2025/11/russian-nuclear-submarine-technology-will-make-north-korean-threat-more-palpable/

Sytas, A., Golubkova, J., Sytas, A., & Golubkova, J. (2026, June 9). Ukraine, Latvia sign drone deal as Russia says Europe not ready to mediate peace talks. *Reuters*. https://www.reuters.com/world/ukraine-latvia-sign-drone-deal-zelenskiy-says-2026-06-09/

The Congressional Research Service. (2024, September 10). *Russia-North Korea Relations*. https://www.everycrsreport.com/reports/IF12760.html

Tiezzi, S. (2021, September 17). *Has the 'China Initiative' Run Its Course?* https://thediplomat.com/2021/09/has-the-china-initiative-run-its-course/

Urbina, F., Lentzos, F., Invernizzi, C., & Ekins, S. (2022). Dual use of artificial-intelligence-powered drug discovery. *Nature Machine Intelligence*, *4*(3), 189–191. https://doi.org/10.1038/s42256-022-00465-9

U.S. Attorney's Office, District of Massachusetts. (2021, January 14). *MIT Professor Arrested and Charged with Grant Fraud*. United States Department of Justice. https://www.justice.gov/usao-ma/pr/mit-professor-arrested-and-charged-grant-fraud

U.S. Attorney's Office, District of Massachusetts. (2022, January 20). *Statement from U.S. Attorney Rachael S. Rollins on the Dismissal of the Gang Chen Case*. United States Department of Justice. https://www.justice.gov/usao-ma/pr/statement-us-attorney-rachael-s-rollins-dismissal-gang-chen-case

Wagner, C. S. (2026). *The Securitization of Science: A Systems Perspective on Policy and Measurement* (arXiv:2605.20313). arXiv. https://doi.org/10.48550/arXiv.2605.20313

Wentzel, J., Graves, Z., & Ball, D. (2026, June 4). *In Support of Mandatory Nucleic Acid Synthesis Screening and Recordkeeping*. The Foundation for American Innovation. https://www.thefai.org/posts/in-support-of-mandatory-nucleic-acid-synthesis-screening-and-recordkeeping

Xie, Y., Lin, X., Li, J., He, Q., & Huang, J. (2023). Caught in the crossfire: Fears of Chinese–American scientists. *Proceedings of the National Academy of Sciences*, *120*(27), e2216248120. https://doi.org/10.1073/pnas.2216248120

Zeidenberg, P. R., & Dearington, M. F. (2022, September 30). *The China Initiative May Have Finally Died—Killed Not by DOJ but the Courts*. ArentFox Schiff. https://www.afslaw.com/perspectives/alerts/the-china-initiative-may-have-finally-died-killed-not-doj-the-courts

Zhang, C. (2025). Market–security tensions in UK–China academic engagement: Perspectives from UK higher education. *Asian Review of Political Economy*, *4*(1), 14. https://doi.org/10.1007/s44216-025-00057-5

ARTICLE 08

From Material Control to Capability Formation: Proliferation-Sensitive Knowledge and UNSCR 1540 Governance

Nanda Kishore

Former Research Fellow, University of Southampton, UK



Abstract

UN Security Council Resolution 1540 (2004) remains central to the prevention of nuclear, chemical and biological weapons proliferation to and by non-State actors. Its practical operation, however, depends on implementation instruments that make some forms of risk more visible than others: domestic laws, export and border controls, matrices, reporting, assistance, outreach and technical reference tools. This article argues that technological change creates a specific visibility problem for that implementation architecture. Proliferation-relevant capability may be assembled through tacit expertise, data, software, digital design information, commercial services and distributed research infrastructures, rather than through the movement of a discrete controlled object. The article develops the concept of proliferation-sensitive knowledge to analyse how such pathways become legible within UNSCR 1540 governance. The term is narrower than security-sensitive knowledge. It refers to knowledge, data, expertise, training, design information or service-based know-how that becomes institutionally connected to WMD-relevant capability formation. Drawing on discursive institutionalism, policy feedback, policy instrumentation and boundary work, the article explains how matrices, reports, assistance mechanisms, outreach, export-control language and voluntary technical reference tools shape what can be

recognised, discussed and supported as 1540-relevant. The resulting framework can structure document analysis, comparison of technology pathways, and awareness, education and research-engagement discussions; it is not a policy roadmap.

Keywords

UNSCR 1540; proliferation-sensitive knowledge; dual-use governance; WMD non-proliferation; implementation architecture; research engagement

Introduction

UN Security Council Resolution 1540 (2004) was designed for a world in which proliferation risk could be made visible through identifiable actors, objects and transfers. Its core concern is the prevention of non-State actors from developing, acquiring, manufacturing, possessing, transporting, transferring or using nuclear, chemical and biological weapons and their means of delivery (United Nations Security Council [UNSC], 2004). Its implementation architecture reflects this concern: national legislation, domestic controls, export and border controls, reporting, assistance mechanisms, matrices, outreach and technical guidance make WMD-related obligations operational across a diverse international system.

Technological change complicates this architecture. Proliferation-relevant capability may now be assembled through less tangible and more distributed pathways: tacit expertise, training, data, software, digital design information, commercial platforms, research collaboration, cloud-based services and distributed innovation infrastructures. These pathways do not always appear as the transfer of a controlled material object. Nor do they map neatly onto listed items, border movements or a single responsible actor. The issue is therefore not simply that knowledge may have harmful uses. That point is already familiar in scholarship on dual-use research, biosecurity, export controls and research security (Atlas & Dando, 2006; Kelle, 2013; Rychnovská, 2016). The more specific issue is how a material-control regime comes to recognise knowledge-based capability formation as a proliferation problem.

This article addresses that issue through the concept of proliferation-sensitive knowledge. The term is deliberately narrower than security-sensitive knowledge. It does not refer to dangerous knowledge in general, and it does not imply a general security claim over science. It refers to knowledge, data, expertise, training, design information or service-based know-how that becomes institutionally connected to WMD-relevant capability formation under UNSCR 1540 governance. The article asks: how does UNSCR 1540's implementation architecture make knowledge-based capability formation visible as a proliferation problem?

The argument is that proliferation-sensitive knowledge is produced through practices of institutional visibility. Matrices, national reports, assistance mechanisms, outreach to academia and industry, export-control guidance, end-use language and voluntary technical reference tools do more than implement a settled mandate. They shape what can be recognised, discussed and supported as relevant to Resolution 1540. This claim draws on four related insights: institutions adapt through discourse and inherited categories (Schmidt, 2008), earlier policy choices structure later possibilities for action (Pierson, 1993, 2000), policy instruments shape what governance systems can classify and address (Lascoumes & Le Galès, 2007), and boundary work helps explain why distinctions between open scientific knowledge, dual-use knowledge and proliferation-sensitive knowledge are produced through institutional practice rather than fixed in advance (Gieryn, 1983).

The article's contribution is not to add another general warning about emerging technologies. It is to explain how an existing material-control regime absorbs knowledge-based capability risks without formally becoming a general technology-control regime. This matters because the visibility threshold is consequential. If the threshold is too restrictive, relevant knowledge-based pathways may remain institutionally invisible. If it is too expansive,

open scientific activity may be unnecessarily securitised. The article therefore shifts attention from the existence of dual-use knowledge to the institutional practices through which knowledge becomes legible as proliferation-relevant.

The analysis proceeds in six sections. The first section characterises UNSCR 1540 as an implementation architecture, not only as a legal text. The second explains technological change as a transformation in capability pathways. The third develops an analytical framework based on discourse, feedback, instruments and boundary work. The fourth defines proliferation-sensitive knowledge as a bounded category of institutional visibility. The fifth uses broad illustrative technology areas rather than a single case. The final section identifies implications for awareness-building, education, research engagement and implementation-support discussion, while avoiding prescriptive policy recommendations to governments.

UNSCR 1540 as an Implementation Architecture

UNSCR 1540 is often described as a binding Security Council resolution on WMD non-proliferation. That description is accurate, but it does not capture the full institutional setting in which the resolution operates. The resolution matters because of the legal obligation it creates, but its day-to-day relevance depends on how implementation is documented, discussed, supported and reviewed. Its architecture includes national reports, 1540 matrices, assistance processes, points of contact, outreach events, technical reference materials and cooperation with international, regional and subregional organisations. The United Nations Security Council 1540 Committee describes the matrix as a primary method for organising information on States' implementation, with information originating mainly from national reports and official government information (United Nations Security Council 1540 Committee, n.d.-a). The matrix is not a neutral mirror of the world. It is a classificatory instrument. It makes certain forms of implementation visible by providing categories into which national measures can be entered.

This matters analytically because implementation tools do more than record compliance. Policy instruments shape how problems are represented. Lascoumes and Le Galès (2007) argue that instruments are not neutral devices; they embody theories of the relationship between governing and governed actors and produce effects of their own. A matrix, an assistance template, a technical reference guide or an outreach category therefore does not merely implement 1540. It structures what can be noticed as a 1540 problem. Legal measures, export controls, border controls, physical protection, end-user controls and enforcement mechanisms are highly visible because they fit the administrative categories of implementation. Diffuse accumulations of knowledge, training, data access and service-based capability are less directly visible until they are translated into those categories.

The original design of UNSCR 1540 also reflects a specific security framing. The resolution responded to concerns that non-State actors could acquire WMD capabilities and that existing treaty regimes did not impose equivalent obligations on all States regarding domestic controls, export controls and non-State actor access. Legal scholarship has debated whether this form of Security Council action represents international legislation and how far such legislative authority can legitimately extend (Asada, 2008; Joyner, 2007; Talmon, 2005). For the purposes of this article, the most important point is not the formal legality debate itself. It is that the resolution's architecture channels technological change through an already established set of categories: State obligation, domestic law, related materials, means of delivery, export and trans-shipment controls, financing, services and enforcement.

International legalisation scholarship helps clarify the institutional form of this arrangement. Resolution 1540 has a binding legal character, yet much of its practical development occurs through softer implementation practices, expert interaction and standardised reporting routines. Abbott et al. (2000) and Abbott and Snidal (2000) show that international governance operates across degrees of obligation, precision and delegation. That perspective is useful here because 1540 combines a hard Security Council mandate with implementation-support practices that depend on information exchange, voluntary technical reference materials and cooperation across organisations. The regime's evolution is therefore not reducible to formal amendment. It is also a process of category formation through implementation.

Successive resolutions have renewed and extended the institutional life of the 1540 Committee, including through comprehensive reviews, assistance facilitation and cooperation with other bodies. Resolution 2663 continued the mandate of the Committee and its Group of Experts and maintained the emphasis on implementation, assistance and cooperation (UNSC, 2022). This continuity is important. The regime has adapted mainly through implementation practice, not through a formal transformation into a new technology-control treaty. Its practical evolution therefore takes place through language, instruments and routines. New technological concerns must be made recognisable inside an architecture built around WMD, non-State actors, related materials and national implementation.

A focus on visibility clarifies this institutional problem. A compliance-centred question asks whether States have fulfilled obligations. A visibility-centred question asks what kinds of proliferation-relevant pathways the implementation architecture is capable of seeing. Compliance remains important, but it is not the whole issue. In a technological environment where capability may emerge through distributed knowledge, services and infrastructures, the capacity to recognise pathways of capability formation becomes analytically prior to the capacity to measure formal implementation. If a risk pathway is not visible within the relevant institutional categories, it may be difficult to discuss through 1540 mechanisms even when it is relevant to WMD non-proliferation awareness.

Technological Change as Capability-Pathway Transformation

The article uses technological change in a deliberately broad sense. It does not focus on one technique or one national context. Emerging technologies are often characterised by radical novelty, relatively fast growth, coherence, prominent impact and uncertainty or ambiguity (Rotolo et al., 2015). Those features matter for 1540 because they alter how capabilities may be produced. The most important transformation is not simply that new tools exist. It is that the relationship between knowledge, materials, services and operational capability becomes less linear. A capability may be assembled by combining open literature, datasets, design software, tacit expertise, outsourced services, commercial platforms and laboratory validation. This mode of assembly is harder to capture through a narrow object-transfer model.

Classic accounts of technology governance emphasise a dilemma of control: early in the development of a technology, consequences are difficult to predict; later, when consequences are clearer, the technology may be harder to redirect (Collingridge, 1980). Anticipatory governance and responsible innovation literatures developed partly as responses to this dilemma, focusing on foresight, reflection, inclusion and responsiveness in the governance of emerging knowledge-based technologies (Guston, 2014; Owen et al., 2012; Stilgoe et al., 2013). This article does not apply those literatures as policy prescriptions. It uses them to clarify why implementation systems face problems of timing and visibility. When capability pathways are still forming, their relevance to 1540 may be uncertain. Once they are fully operationalised, they may already sit across multiple institutions, markets and infrastructures.

The concept of knowledge-based capability formation captures this shift. It refers to processes through which knowledge, data, skill, technical routines, services and infrastructures combine to generate operational capacity. In the proliferation context, capability formation is more specific: it concerns pathways that may become connected to nuclear, chemical or biological weapons, their means of delivery, related materials or non-State actor acquisition. The term avoids two mistakes. It avoids reducing risk to discrete material objects. It also avoids treating knowledge as intrinsically sensitive. The key question is relational: when does a knowledge pathway become linked to WMD-relevant capability formation in a way that UNSCR 1540 governance can recognise?

This shift also matters for non-proliferation scholarship on dual-use trade. Fuhrmann (2008) shows that dual-use commerce sits at the intersection of economic exchange and security concern. The present article extends that insight from dual-use trade in items to knowledge-based capability formation. The relevant pathway may include

trade, but it may also include technical training, data access, design workflows, service provision or collaboration. This is why proliferation sensitivity cannot be inferred from the existence of technical knowledge alone. It has to be analysed through the relation between knowledge, capability, access and institutional uptake.

The life sciences illustrate this transformation with particular clarity. The synthesis of poliovirus from cDNA showed that sequence information could be connected to the generation of infectious virus in the absence of a natural template (Cello et al., 2002). The reconstruction of the 1918 influenza virus triggered debate about the publication of dual-use information, scientific value and biosecurity (Tumpey et al., 2005). The later construction of an infectious horsepox virus from chemically synthesised DNA fragments revived questions about gene synthesis, know-how and access to enabling services (Noyce et al., 2018). These examples do not imply that life-science publication is generally proliferation-sensitive. They show how knowledge, sequence information, laboratory skill and commercial synthesis can combine in ways that make capability formation a governance concern.

Similar patterns appear in other technology areas. AI-enabled drug discovery has demonstrated that models and data used for beneficial pharmaceutical research can also be redirected toward the generation of toxic molecules (Urbina et al., 2022). Gene drives and synthetic biology raise questions about design, release, ecological effects and distributed research communities (Kelle, 2013; Oye et al., 2014). Digital manufacturing and cloud-based technical services can separate design knowledge, production capacity and physical location. These domains differ technically, but they share a governance feature: capability can be distributed across knowledge, software, data, service access and tacit expertise rather than concentrated in one controlled item.

The concept of proliferation-sensitive knowledge therefore does not designate a fixed list of fields. It identifies a relation between knowledge and capability formation. A textbook chapter, a published protocol, a dataset, a design file, a training session or a commercial service may remain ordinary scientific or technical activity. It becomes proliferation-sensitive only when it is institutionally linked to WMD-relevant capability formation, control circumvention, means of delivery, related materials or non-State actor access. This threshold is crucial. Potential harm alone is too broad. Many forms of engineering, chemistry, biology, computing and materials science can cause harm under some conditions. Proliferation sensitivity requires a narrower connection to the 1540 problem space.

Analytical Framework: Discourse, Feedback, Instruments and Boundaries

The framework is organised as a sequence of four analytical questions rather than as a catalogue of theories. Discursive institutionalism asks how a technological development is described in 1540 language; policy feedback asks how the resolution's inherited WMD and non-State actor framing limits that description; policy instrumentation asks where the description is stabilised in matrices, reports, guidance or assistance practices; and boundary work asks how institutions distinguish open science, dual-use knowledge and proliferation-sensitive knowledge. Used together, the four tools connect interpretation, institutional inheritance, administrative practice and threshold-setting (Schmidt, 2008; Pierson, 1993; Lascoumes & Le Galès, 2007; Gieryn, 1983).

Discursive institutionalism identifies the vocabulary through which unfamiliar technologies become intelligible to the regime. A cloud laboratory, gene synthesis service or machine-learning workflow is not self-evidently a 1540 object; actors must describe it as a route to capability formation, a service, an intangible transfer, an end-use concern or an awareness-building topic. In practical analysis, this directs attention to the wording used in national reports, Committee materials, outreach programmes and expert discussions: which labels are used, which WMD connections are asserted, and which features remain unnamed (Schmidt, 2008).

Policy feedback explains why those descriptions remain anchored in the resolution's original architecture. The 2004 focus on WMD, means of delivery, related materials, non-State actors and State obligations supplies an interpretive template for later technological concerns. A practical application is to compare whether an emerging

pathway is linked to those established categories or discussed only as a general technology risk. The comparison shows where inherited categories enable recognition and where they filter a pathway out of 1540-specific discussion (Pierson, 1993, 2000).

Policy instrumentation locates institutional translation in observable tools. Matrices, reports, assistance mechanisms and voluntary technical reference materials classify implementation and make particular forms of evidence portable across States and organisations. The 1540 Committee's export-control reference guide, for example, includes catch-all controls, end-user controls, intangible transfers, export-related services and technical assistance (United Nations Security Council 1540 Committee, n.d.-b). These categories provide practical coding points for analysing where knowledge and service-based pathways enter the regime, which instruments carry them, and which pathways remain difficult to record.

Boundary work provides the limiting test. The distinction between open scientific knowledge, dual-use knowledge and proliferation-sensitive knowledge is produced through classification, review, outreach and professional practice, not discovered as a fixed technical property (Gieryn, 1983; Jasanoff, 2004, 2005). In practical terms, boundary analysis asks what evidence is used to justify movement from one category to another and whether that evidence includes WMD connection, operational specificity, access and institutional uptake. This application also incorporates the securitisation warning: a 1540-specific category should diagnose bounded capability pathways rather than convert advanced science into a general object of suspicion (Buzan et al., 1998).

Producing Proliferation-Sensitive Knowledge

Proliferation-sensitive knowledge is produced when knowledge-based activity becomes visible to 1540 governance as connected to WMD-relevant capability formation. Production here means institutional classification: a pathway is narrated, reported, reviewed and connected to implementation categories. The framework operationalises this process through four visibility tests—capability proximity, operational specificity, access pathway and institutional uptake. These tests can structure document analysis, comparison of technology cases and awareness or research-engagement exercises; they are analytical prompts, not legal criteria or a control checklist.

Capability proximity and operational specificity assess the substantive connection between knowledge and capability. The first asks whether a source identifies a concrete link to nuclear, chemical or biological weapons, means of delivery, related materials or supporting systems. The second asks whether the knowledge supplies actionable parameters, validated procedures, designs, troubleshooting or tacit know-how. Applied to a report, publication or technology case, the two tests distinguish general education and broad scientific principles from information that materially shortens the path to a WMD-relevant capability (Miller & Selgelid, 2007; Patrone et al., 2012; Selgelid, 2009).

Access pathway and institutional uptake assess movement and recognition. The access-pathway test maps how knowledge travels through publications, training, datasets, software repositories, commercial services, cloud laboratories, expert networks or collaboration. The institutional-uptake test then checks whether that pathway appears in matrices, national reports, guidance, assistance discussions, outreach or expert discourse. Used together, they allow an analyst to trace distributed capability assembly and to identify the point at which a technical possibility becomes institutionally legible within 1540 governance.

Table 1. Visibility tests and practical analytical applications

Visibility test	Core question	Boundary clarified	Practical analytical use
Capability proximity	How close is the knowledge to WMD-relevant capability formation?	General education from capability-enabling knowledge	Code the specific WMD-related function or enabling capability
Operational specificity	Does it provide actionable procedures, parameters, designs or troubleshooting?	General principles from operational know-how	Distinguish broad principles from actionable workflows or tacit know-how

Access pathway	How does the knowledge travel through training, service, platform, publication or collaboration?	A discrete transfer from distributed capability assembly	Map the actors, platforms and services through which capability is assembled
Institutional uptake	Has the pathway entered matrices, reports, guidance, outreach or expert discourse?	Technical possibility from institutionally recognised 1540 relevance	Locate the pathway in reports, matrices, guidance, assistance or outreach

Source: Author, drawing on the analytical framework developed in this article.

Used comparatively, the tests help avoid two errors. Under-recognition occurs when analysis looks only for controlled materials or equipment and misses data, tacit skills, service access or platform-mediated capability. Over-recognition occurs when all advanced knowledge is treated as proliferation-sensitive because it has possible harmful uses. The narrower threshold requires a demonstrable connection to WMD-relevant capability formation and an account of how that connection becomes articulable within 1540's implementation architecture.

This also clarifies the relation between open knowledge, dual-use knowledge and proliferation-sensitive knowledge. Open scientific knowledge refers to knowledge circulated for education, research and legitimate public benefit. Dual-use knowledge refers to knowledge with potential beneficial and harmful applications. Proliferation-sensitive knowledge is a subset: knowledge that becomes institutionally linked to WMD-relevant capability formation under 1540 governance. The category is therefore relational and institutional. It is not a claim about the moral character of a field, researcher or technology.

The point can be made through a simple example. Learning digital manufacturing is not proliferation-sensitive merely because objects produced by digital manufacturing can cause harm. General technical education remains outside the core 1540 concern. It becomes relevant only if knowledge, design files, production services or tacit manufacturing expertise connect to WMD-related components, means of delivery, controlled equipment or control circumvention. The same logic applies to life-science techniques, AI models or chemical design tools. Potential harm is not enough. Proliferation sensitivity requires a structured connection to the regime's object: preventing WMD-related capabilities from being acquired or used by non-State actors.

Illustrative Technology Areas

This section uses broad illustrative technology areas to show how the argument works. The aim is not to select one technology as the object of control. It is to demonstrate how different fields expose the same visibility problem: proliferation-relevant capability may be distributed across knowledge, data, services, expertise and infrastructure.

Life sciences and synthetic biology

Life sciences and synthetic biology have long occupied a central place in dual-use debate. The National Research Council's Fink report framed the destructive application of biotechnology as a governance problem requiring awareness, education, review and scientific responsibility (National Research Council, 2004). A later joint report by the Institute of Medicine and National Research Council situated life-science developments within broader international networks of research, industry and public health (Institute of Medicine & National Research Council,

2006). Atlas and Dando (2006) similarly argued that life-science knowledge can be open and beneficial while also requiring a culture of responsible conduct. These debates are directly relevant to 1540 because biological weapons, related materials and non-State actor access sit within the resolution's core scope.

Synthetic biology sharpens the issue because it shifts attention from possession of an organism to the capacity to design, order, assemble, modify or validate biological systems. Carlson (2009) documented the changing economics of DNA synthesis, indicating how technological capability could become more accessible over time. Kelle (2013) analysed synthetic biology as a dual-use governance problem and criticised patchwork approaches. Kobokovich et al. (2019) examined gene synthesis governance and screening as a security issue. These sources show that capability formation in the life sciences can involve sequence data, synthesis services, customer screening, laboratory tacit knowledge and publication practices. That mix is difficult to reduce to a single controlled object.

Recent assessments of biotechnology and the bioeconomy reinforce this point. The National Academies of Sciences, Engineering, and Medicine (2018, 2020) frame synthetic biology and the bioeconomy as domains in which beneficial innovation, data systems, biological design capacity and security awareness develop together. For the purposes of this article, the importance of this literature is not that biotechnology should be treated as a single priority technology. It is that life-science capability increasingly depends on integrated knowledge systems: sequence information, design tools, synthesis capacity, tacit laboratory skill, quality assurance and service infrastructures.

The life-science examples also demonstrate why boundary discipline matters. Cello et al. (2002), Tumpey et al. (2005) and Noyce et al. (2018) are frequently discussed because they connect published sequence or method information to the reconstruction or synthesis of viruses. Yet these examples do not justify treating all virology or synthetic biology as proliferation-sensitive. They show that certain configurations of sequence information, operational know-how, commercial synthesis and validation capacity can become connected to biological capability formation. The object of analysis is the pathway, not the field as a whole.

AI, data systems and computational design

AI and data systems represent a second broad area. AI can support beneficial research by accelerating discovery, modelling, design and screening. It can also alter the pathway from knowledge to capability by lowering the burden of search, optimisation and design. Urbina et al. (2022) showed how an AI drug-discovery system could be redirected toward toxic molecule generation. The relevance for this article is not a general claim that AI is a proliferation problem. The relevance is that data, models and computational workflows can function as capability accelerators when connected to chemical or biological design tasks.

This creates a visibility challenge for 1540 because AI systems may not be weapons, controlled materials or export items in the conventional sense. Their significance may lie in how they combine with domain data, laboratory services, synthesis providers and user expertise. The proliferation-sensitive element may be located not in the model alone but in the configuration: a model, a dataset, prompts or workflows, target knowledge, validation access and downstream synthesis or testing. The category of operational specificity is useful here. A general AI system is not proliferation-sensitive as such. A validated workflow that substantially reduces barriers to WMD-relevant capability formation is closer to the relevant threshold.

Information-hazard debates sharpen the same point. Esvelt (2018) argues that scientific communication can sometimes disseminate true information that may enable harmful action under particular conditions. That argument does not justify treating all publication as a proliferation issue. It does, however, support the narrower claim developed here: some knowledge pathways become relevant when operational specificity, access and capability proximity combine in ways that can be translated into WMD-relevant institutional categories.

The AI example also demonstrates the importance of institutional uptake. If a model-based capability pathway is framed only as an abstract emerging-technology concern, it may remain outside 1540 discussion. If it is translated into technical assistance, end-use review, research-security education or awareness-building for relevant scientific communities, it becomes visible within implementation-support practice. This translation is not automatic. It depends on expert discourse, institutional categories and the fit between technological features and 1540's WMD mandate.

Digital manufacturing, services and distributed infrastructures

Digital manufacturing, platform services and distributed infrastructures represent a third broad area. They are useful because they show the weakness of potential harm as a criterion. A technology such as additive manufacturing may be used for ordinary prototyping, education, medical devices or industrial production. It may also be used to produce components with security relevance. The difference is not inherent in the technology itself. It lies in object connection, capability proximity, end-use context, design specificity and institutional framing.

Service infrastructures make the issue sharper. A user may access laboratory capacity, computing power, synthesis providers, modelling tools or manufacturing services without owning the underlying equipment. Capability formation may occur through orchestration rather than possession. This does not make services generally proliferation-sensitive. It means that services can become part of a proliferation-relevant pathway when they provide access to WMD-related capability formation, controlled components, tacit operational support or control circumvention. The United Nations Security Council 1540 Committee's export-control reference language already contains entry points for services, end-user controls, intangible transfer and technical assistance (United Nations Security Council 1540 Committee, n.d.-b). The analytical issue is how those entry points capture distributed capability formation.

The same logic can be seen in responsible innovation and anticipatory governance scholarship. Collingridge (1980) identified the difficulty of steering technologies when consequences are uncertain and social embedding increases over time. Guston (2014), Owen et al. (2012) and Stilgoe et al. (2013) later developed approaches that emphasise anticipation, reflexivity, inclusion and responsiveness. This article does not import those literatures as a policy prescription. It uses them to frame why early awareness and research engagement matter before knowledge-based pathways are fully stabilised as risks.

These examples confirm the article's central point. Technological change does not simply add new items to a list. It changes the forms through which capability is assembled. In some cases, capability formation remains visible because it involves controlled materials or equipment. In others, the relevant pathway is knowledge-intensive, service-based or platform-mediated. 1540 governance can address only what it can classify through its mandate and instruments. Visibility is therefore a core problem, not a secondary administrative detail.

Visibility Gaps and Boundary Discipline

A visibility approach identifies four gaps in the 1540 implementation architecture. These are not accusations of failure. They are analytical distinctions that help locate where knowledge-based capability formation may be difficult to recognise.

First, there is an object gap. The resolution is highly effective at naming weapons, means of delivery and related materials, and implementation tools are well suited to documenting legal and regulatory controls around identifiable items. Yet knowledge-based capability formation may not appear as an object. It may take the form of data, expertise, design logic, platform access or workflow integration. This creates a gap between object-based visibility and capability-based reality.

Second, there is a transfer gap. Traditional control language often assumes a transfer that can be identified, licensed, refused or penalised. Knowledge may travel through less discrete routes: remote collaboration, cloud access, training, publications, code repositories, technical troubleshooting or service provision. Some of these routes can be captured through intangible transfer or technical assistance categories, but not all cumulative knowledge flows appear as a single regulated transfer.

Third, there is an actor gap. UNSCR 1540 places legal obligations on States, yet knowledge-based capability formation often occurs through universities, firms, journals, funders, professional societies, service providers, technical experts and international research networks. This does not displace the State as the legal duty-bearer. It does mean that awareness and implementation-support discussions increasingly pass through non-State knowledge infrastructures. Rychnovská (2016) shows how responsible science discourses can ethicalise security by assigning responsibility to scientific communities. The 1540 context adds another layer: scientific communities become relevant not because they are legally equivalent to States, but because their knowledge infrastructures can shape visibility and awareness.

Fourth, there is a threshold gap. It is difficult to specify when ordinary knowledge becomes dual-use, and when dual-use knowledge becomes proliferation-sensitive. The literature on dual-use research has wrestled with this problem for two decades (Kuhlau et al., 2008; Miller & Selgelid, 2007; Patrone et al., 2012; Selgelid, 2009). A 1540-specific threshold must be narrower than general dual-use concern. It requires connection to WMD-relevant capability formation and institutional translation into the 1540 mandate. Without this discipline, the concept risks becoming too broad. With too strict a threshold, knowledge-based pathways may remain invisible until they are already operationalised.

The framework in this article therefore aims to make a specific contribution: it shifts discussion from 'what knowledge is dangerous?' to 'how does 1540 governance make knowledge-based capability formation visible?' The first question is too broad and invites over-securitisation. The latter question is institutionally bounded. It asks how categories, instruments and discourses shape what the regime can see, and what kinds of awareness or research engagement can be supported within a non-prescriptive academic framing.

Boundary discipline also protects open science. Open publication, education and international collaboration are central to scientific progress. The 1540 framework does not become more coherent by treating openness itself as a problem. The analytical task is to identify conditions under which knowledge pathways move closer to proliferation-relevant capability formation. Those conditions include operational specificity, direct WMD connection, service-enabled access, control circumvention and institutional uptake. The more clearly the threshold is defined, the less likely the concept of proliferation-sensitive knowledge will become an unlimited claim over scientific knowledge.

Relevance for Awareness, Education and Research Engagement

The purpose of this article is academic analysis, not a policy roadmap. Its practical value lies in providing a common method for awareness-building, education, research engagement and implementation-support discussion. The four visibility tests translate the abstract claim about institutional visibility into questions that can be applied to documents, technology pathways and interdisciplinary dialogue while keeping the analysis bounded by UNSCR 1540's WMD mandate.

For awareness-building, the framework can be applied to short scenarios. Participants can compare, for example, general technical education with a validated workflow that combines design information, service access and WMD-relevant end use. Asking the four visibility questions makes the reason for concern explicit and prevents 'security' from operating as an undifferentiated claim against academic openness. It also helps non-proliferation audiences move from item lists toward capability pathways without losing the resolution's scope.

For education, the tests provide a structured exercise for comparing publications, datasets, training activities or commercial services. Students can identify the relevant capability connection, assess the level of operational detail, map the access pathway and determine whether the issue has entered 1540-related instruments or discourse. The exercise demonstrates why the same technique may remain ordinary scientific activity in one configuration and become proliferation-sensitive in another.

For research engagement, the framework can organise dialogue between technical and non-proliferation communities. Technical experts can specify feasibility, validation requirements, tacit constraints and legitimate uses; non-proliferation experts can map actors, end uses, control circumvention and available implementation categories. The shared output is a pathway analysis showing where knowledge moves closer to WMD-relevant capability formation and where institutional vocabulary does or does not capture that movement.

For implementation-support discussion, the framework can be used to read matrices, national reports, assistance requests, outreach materials and voluntary reference tools as infrastructures of visibility. A comparative review can code which knowledge pathways are represented through categories such as intangible transfer, technical assistance, services or end use, and where object, transfer, actor or threshold gaps persist. This use identifies analytical blind spots and discussion needs without prescribing new legal obligations or national controls.

The article also helps explain why the choice of terminology matters. 'Security-sensitive knowledge' is too open-ended for this project because it does not specify the relevant security domain. 'Proliferation-sensitive knowledge' is more precise. It directs attention to WMD-relevant capability formation, non-State actor access and the institutional logic of UNSCR 1540. This terminological narrowing is not a cosmetic change. It is what prevents the article from drifting into generalised securitisation of science. It keeps the analysis anchored in the resolution's scope.

Positioning the Contribution in Existing Literatures

The contribution of the article becomes clearer when placed against three literatures. The first is the literature on UNSCR 1540 and Security Council legislative authority. This literature asks how Resolution 1540 fits within the legal powers of the Security Council, whether it represents an unusual form of global legislation and how its legitimacy can be assessed (Asada, 2008; Joyner, 2007; Talmon, 2005). That debate is essential background, but it does not fully explain how technological change is absorbed through implementation practice. A legal-authority approach can tell us why the resolution is binding and why its scope matters. It does not by itself explain how matrices, reference guides, outreach activities and assistance processes classify new knowledge-based pathways. This article therefore treats the legal mandate as a starting condition and asks how that mandate is operationalised through visibility practices.

The second literature is the implementation and compliance literature on 1540. Crail (2006) provides an early risk-based approach to implementation, while Holmes and von Winterfeldt (2025) develop a more recent compliance evaluation method. These works are valuable because they take implementation seriously rather than treating the resolution as only a formal legal instrument. The present article builds on that orientation but shifts the unit of analysis. It is less concerned with whether a State has implemented a measure and more concerned with what kinds of technological and knowledge pathways the implementation architecture can recognise. This is a movement from compliance assessment to visibility analysis. The shift matters because a State may have formal controls while certain knowledge-based capability pathways remain difficult to observe, classify or discuss through available categories.

The third literature is the literature on dual-use research and responsible science. This literature has already made the strongest version of the point that knowledge, publication, training and scientific openness can raise security questions (Atlas & Dando, 2006; Kuhlau et al., 2008; Miller & Selgelid, 2007; Patrone et al., 2012; Rychnovská, 2016; Selgelid, 2009). The article does not repeat that argument. Its contribution is to ask what happens when dual-use knowledge debates enter a specific non-proliferation implementation architecture. The logic of 1540 is not

identical to general biosecurity ethics or responsible innovation. It is narrower, because it is anchored in WMD, means of delivery, related materials and non-State actor access. It is also more institutional, because it operates through reports, matrices, assistance and cooperation mechanisms. Proliferation-sensitive knowledge therefore has to be understood differently from dual-use knowledge in general.

The central intervention is therefore to study recognition rather than restriction. Seen in this way, proliferation-sensitive knowledge is a lens for institutional analysis rather than a label for prohibition. It helps identify when a knowledge pathway has moved close enough to WMD-relevant capability formation to become meaningful within the 1540 vocabulary. The question is not whether a field is dangerous in the abstract. The question is how knowledge, access, services and expertise are assembled, interpreted and made visible through an existing non-proliferation regime.

The article also avoids treating implementation instruments as merely technical paperwork. Their significance is epistemic as well as administrative. They stabilise categories, create common vocabularies and make certain forms of evidence portable across States, experts and organisations. This is why a matrix or reference guide can matter for knowledge governance even when it does not create a new legal obligation. It shapes the field of possible recognition.

This positioning clarifies the analytical significance of the article. If the conclusion were merely that emerging technologies create new knowledge risks, the argument would be weak. The stronger claim is that technological change creates classification pressure inside an existing regime. The regime must maintain a bounded WMD mandate, but it also has to recognise capability pathways that do not always look like the movement of materials or equipment. This is a problem of institutional translation. Knowledge-based pathways become relevant to 1540 only after they are translated into categories such as technical assistance, services, intangible transfer, research engagement, end-use concern or implementation support. The article's object is that translation process.

The concept of proliferation-sensitive knowledge is therefore not offered as a new label for already familiar risk. It is offered as a way to describe the institutional result of translation. A form of knowledge may begin as open science, become discussable as dual-use knowledge and become proliferation-sensitive only when it is connected to WMD-relevant capability formation through recognised institutional pathways. This sequential framing matters because it avoids both determinism and alarmism. It does not assume that technological capability mechanically produces governance expansion. It also does not assume that governance must expand wherever risk can be imagined. It studies how particular categories make expansion possible without turning 1540 into a general technology-control regime.

The distinction between visibility and prescription is especially important for this project. A prescriptive article would move from identifying a gap to telling States or institutions how to close it. That is not the move made here. The article instead explains why certain gaps arise and why they matter for shared understanding. It maps the categories through which knowledge-based capability formation becomes intelligible. That mapping is analytically useful even without a policy roadmap because it clarifies the assumptions built into existing instruments, the limits of item-centred thinking and the conditions under which research communities may become relevant to 1540 awareness. In that sense, the practical value of the article lies in conceptual discipline: it gives readers a way to discuss sensitive knowledge pathways without collapsing all advanced research into proliferation risk.

This argument has consequences for how examples are used. A single technology case can easily distort the analysis. If the article focused only on gene synthesis, it might become a life-science biosecurity article. If it focused only on 3D printing, it might drift toward conventional weapons or export-control debates outside the core WMD frame. If it focused only on AI, it might become a general AI governance article. Broad illustrative examples are therefore more appropriate. They show that the central phenomenon is not one technology. It is the recurrence of a capability-pathway problem across different technological fields. Knowledge, data, service access and tacit expertise become relevant when they combine with WMD-relevant capability formation.

The article also contributes a vocabulary for analysing the limits of visibility. It identifies object, transfer, actor and threshold gaps. Object gaps arise when risk is not attached to a discrete material item. Transfer gaps arise when capability formation occurs through cumulative or intangible routes rather than a single shipment. Actor gaps arise when legally State-centred obligations depend on universities, firms, technical experts, platforms and professional communities for awareness and implementation support. Threshold gaps arise when the boundary between open science, dual-use knowledge and proliferation-sensitive knowledge is difficult to define. These gaps are not defects that can be solved by one policy instruction. They are structural features of governing knowledge-based capability formation through a regime built around material control.

A further benefit of the visibility approach is that it gives the article a neutral global frame. The analysis does not identify one State as the problem, one government as the target audience or one national system as the model. It examines the architecture of a Security Council regime and the broader research and technology environments in which it operates. This is consistent with the publication's emphasis on academic perspectives, awareness-building, education, research engagement, good practices and implementation-support discussions. The article can therefore speak to researchers, technical communities, international organisations and non-proliferation audiences without becoming a national policy assessment.

Conclusion

UNSCR 1540 remains a material-control regime in its legal architecture, but technological change increasingly exposes knowledge-based pathways of capability formation. This article has argued that the resulting problem is best understood as a visibility problem. The issue is not that the world has only recently discovered that knowledge can be misused. The issue is how an implementation architecture designed around State obligations, domestic controls, export controls, reporting, assistance and outreach makes some forms of knowledge-based capability formation visible as proliferation problems.

The article has defined proliferation-sensitive knowledge as knowledge, data, expertise, training, design information or service-based know-how that becomes institutionally linked to WMD-relevant capability formation under UNSCR 1540 governance. This definition is intentionally bounded. It does not apply to advanced knowledge in general. It requires a connection to WMD, means of delivery, related materials, non-State actor access, control circumvention or implementation categories. The category is therefore relational, institutional and regime-specific.

The analytical framework combines discourse, feedback, instruments and boundary work. Discursive institutionalism explains how technological change is translated into 1540 language; policy feedback identifies the inherited categories that condition that translation; policy instrumentation locates it in matrices, reports, assistance processes and technical reference tools; and boundary work disciplines the distinction between open science, dual-use knowledge and proliferation-sensitive knowledge. Operationalised through capability proximity, operational specificity, access pathway and institutional uptake, the framework can be used to compare documents and technology pathways while preserving the resolution's bounded mandate.

The article's main contribution is to move beyond the obvious claim that knowledge can be risky. It shows that visibility is a central feature of 1540 governance. What the regime can see is shaped by its instruments; what it cannot easily see may remain outside structured discussion. This insight matters for awareness, education and research engagement because it creates a disciplined way to talk about emerging technologies without treating scientific openness as a general threat. It also helps scholars analyse how older non-proliferation regimes adapt to technological change through categories, instruments and institutional practice.

Proliferation-sensitive knowledge is therefore not the starting point of the article. It is the outcome to be explained. It is produced when knowledge-based capability formation becomes visible through the institutional practices of UNSCR 1540 governance. That is the article's central claim and its contribution to the TechHorizon publication project.

References

Abbott, K. W., Keohane, R. O., Moravcsik, A., Slaughter, A.-M., & Snidal, D. (2000). The concept of legalization. *International Organization*, 54(3), 401-419. https://doi.org/10.1162/002081800551271

Abbott, K. W., & Snidal, D. (2000). Hard and soft law in international governance. *International Organization*, 54(3), 421-456. https://doi.org/10.1162/002081800551280

Asada, M. (2008). Security Council Resolution 1540 to combat WMD terrorism: Effectiveness and legitimacy in international legislation. *Journal of Conflict and Security Law*, 13(3), 303-332. https://doi.org/10.1093/jcsl/krp002

Atlas, R. M., & Dando, M. (2006). The dual-use dilemma for the life sciences: Perspectives, conundrums, and global solutions. *Biosecurity and Bioterrorism: Biodefense Strategy, Practice, and Science*, 4(3), 276-286. https://doi.org/10.1089/bsp.2006.4.276

Buzan, B., Wæver, O., & de Wilde, J. (1998). *Security: A new framework for analysis.* Lynne Rienner Publishers.

Carlson, R. (2009). The changing economics of DNA synthesis. *Nature Biotechnology*, 27(12), 1091-1094. https://doi.org/10.1038/nbt1209-1091

Cello, J., Paul, A. V., & Wimmer, E. (2002). Chemical synthesis of poliovirus cDNA: Generation of infectious virus in the absence of natural template. *Science*, 297(5583), 1016-1018. https://doi.org/10.1126/science.1072266

Collingridge, D. (1980). *The social control of technology*. Frances Pinter.

Crail, P. (2006). Implementing UN Security Council Resolution 1540: A risk-based approach. *The Nonproliferation Review*, 13(2), 355-399. https://doi.org/10.1080/10736700601012193

Esvelt, K. M. (2018). Inoculating science against potential pandemics and information hazards. *PLOS Pathogens*, 14(10), e1007286. https://doi.org/10.1371/journal.ppat.1007286

Fuhrmann, M. (2008). Exporting mass destruction? The determinants of dual-use trade. *Journal of Peace Research*, 45(5), 633-652. https://doi.org/10.1177/0022343308094324

Gieryn, T. F. (1983). Boundary-work and the demarcation of science from non-science: Strains and interests in professional ideologies of scientists. *American Sociological Review*, 48(6), 781-795. https://doi.org/10.2307/2095325

Guston, D. H. (2014). Understanding anticipatory governance. *Social Studies of Science*, 44(2), 218-242. https://doi.org/10.1177/0306312713508669

Holmes, J., & von Winterfeldt, D. (2025). Evaluating compliance with UN Security Council Resolution 1540 on nonproliferation of weapons of mass destruction. *Risk Analysis*, 45(7), 1766-1791. https://doi.org/10.1111/risa.17697

Institute of Medicine, & National Research Council. (2006). Globalization, biosecurity, and the future of the life sciences. *National Academies Press*. https://doi.org/10.17226/11567

Jasanoff, S. (Ed.). (2004). *States of knowledge: The co-production of science and the social order*. Routledge.

Jasanoff, S. (2005). *Designs on nature: Science and democracy in Europe and the United States.* Princeton University Press.

Joyner, D. H. (2007). Non-proliferation law and the United Nations system: Resolution 1540 and the limits of the power of the Security Council. *Leiden Journal of International Law*, 20(2), 489-518. https://doi.org/10.1017/S0922156507004153

Kelle, A. (2013). Beyond patchwork precaution in the dual-use governance of synthetic biology. *Science and Engineering Ethics*, 19(3), 1121-1139. https://doi.org/10.1007/s11948-012-9365-8

Kobokovich, A., West, R., Montague, M., Inglesby, T., & Gronvall, G. K. (2019). Strengthening security for gene synthesis: Recommendations for governance. *Health Security*, 17(6), 419-429. https://doi.org/10.1089/hs.2019.0110

Kuhlau, F., Eriksson, S., Evers, K., & Höglund, A. T. (2008). Taking due care: Moral obligations in dual use research. *Bioethics*, 22(9), 477-487. https://doi.org/10.1111/j.1467-8519.2008.00695.x

Lascoumes, P., & Le Galès, P. (2007). Introduction: Understanding public policy through its instruments - From the nature of instruments to the sociology of public policy instrumentation. *Governance*, 20(1), 1-21. https://doi.org/10.1111/j.1468-0491.2007.00342.x

Miller, S., & Selgelid, M. J. (2007). Ethical and philosophical consideration of the dual-use dilemma in the biological sciences. *Science and Engineering Ethics*, 13(4), 523-580. https://doi.org/10.1007/s11948-007-9043-4

National Academies of Sciences, Engineering, and Medicine. (2018). *Biodefense in the age of synthetic biology*. National Academies Press. https://doi.org/10.17226/24890

National Academies of Sciences, Engineering, and Medicine. (2020). *Safeguarding the bioeconomy*. National Academies Press. https://doi.org/10.17226/25525

National Research Council. (2004). *Biotechnology research in an age of terrorism.* National Academies Press. https://doi.org/10.17226/10827

Noyce, R. S., Lederman, S., & Evans, D. H. (2018). Construction of an infectious horsepox virus vaccine from chemically synthesized DNA fragments. *PLOS ONE*, 13(1), e0188453. https://doi.org/10.1371/journal.pone.0188453

Owen, R., Macnaghten, P., & Stilgoe, J. (2012). Responsible research and innovation: From science in society to science for society, with society. *Science and Public Policy*, 39(6), 751-760. https://doi.org/10.1093/scipol/scs093

Oye, K. A., Esvelt, K., Appleton, E., Catteruccia, F., Church, G., Kuiken, T., Lightfoot, S. B.-Y., McNamara, J., Smidler, A., & Collins, J. P. (2014). Regulating gene drives. *Science*, 345(6197), 626-628. https://doi.org/10.1126/science.1254287

Patrone, D., Resnik, D., & Chin, L. (2012). Biosecurity and the review and publication of dual-use research of concern. *Biosecurity and Bioterrorism: Biodefense Strategy, Practice, and Science*, 10(3), 290-298. https://doi.org/10.1089/bsp.2012.0011

Pierson, P. (1993). When effect becomes cause: Policy feedback and political change. *World Politics*, 45(4), 595-628. https://doi.org/10.2307/2950710

Pierson, P. (2000). Increasing returns, path dependence, and the study of politics. *American Political Science Review*, 94(2), 251-267. https://doi.org/10.2307/2586011

Rotolo, D., Hicks, D., & Martin, B. R. (2015). What is an emerging technology? *Research Policy*, 44(10), 1827-1843. https://doi.org/10.1016/j.respol.2015.06.006

Rychnovská, D. (2016). Governing dual-use knowledge: From the politics of responsible science to the ethicalization of security. *Security Dialogue*, 47(4), 310-328. https://doi.org/10.1177/0967010616658848

Schmidt, V. A. (2008). Discursive institutionalism: The explanatory power of ideas and discourse. *Annual Review of Political Science*, 11, 303-326. https://doi.org/10.1146/annurev.polisci.11.060606.135342

Selgelid, M. J. (2009). Governance of dual-use research: An ethical dilemma. *Bulletin of the World Health Organization*, 87(9), 720-723. https://doi.org/10.2471/BLT.08.051383

Stilgoe, J., Owen, R., & Macnaghten, P. (2013). Developing a framework for responsible innovation. *Research Policy*, 42(9), 1568-1580. https://doi.org/10.1016/j.respol.2013.05.008

Talmon, S. (2005). The Security Council as world legislature. *American Journal of International Law*, 99(1), 175-193. https://doi.org/10.2307/3246097

Tumpey, T. M., Basler, C. F., Aguilar, P. V., Zeng, H., Solórzano, A., Swayne, D. E., Cox, N. J., Katz, J. M., Taubenberger, J. K., Palese, P., & García-Sastre, A. (2005). Characterization of the reconstructed 1918 Spanish influenza pandemic virus. *Science*, 310(5745), 77-80. https://doi.org/10.1126/science.1119392

United Nations Security Council 1540 Committee. (n.d.-a). *1540 matrices*. United Nations. https://www.un.org/en/sc/1540/national-implementation/1540-matrices.shtml

United Nations Security Council 1540 Committee. (n.d.-b). *Voluntary Technical Reference Guide - Export Controls*. United Nations. https://www.un.org/en/sc/1540/national-implementation/vTRG-ExpCtrl-1.shtml

United Nations Security Council. (2004). Resolution 1540 (2004), S/RES/1540. https://digitallibrary.un.org/record/520326

United Nations Security Council. (2022). Resolution 2663 (2022), S/RES/2663. https://digitallibrary.un.org/record/3996316

Urbina, F., Lentzos, F., Invernizzi, C., & Ekins, S. (2022). Dual use of artificial-intelligence-powered drug discovery. *Nature Machine Intelligence*, 4(3), 189-191. https://doi.org/10.1038/s42256-022-00465-9

ARTICLE 09

From Classroom to Commercialisation: Biomedical Engineering Capability Formation and UNSCR 1540 Governance

Spyros Kamins

PhD from Aston University



Abstract

This article examines how United Nations Security Council Resolution (UNSCR) 1540 should address capability formation across biomedical engineering education and innovation ecosystems. While UNSCR 1540 was originally designed to prevent non-State actors from acquiring weapons of mass destruction (WMD), related materials and means of delivery, contemporary biomedical innovation ecosystems increasingly generate capabilities through education, doctoral training, laboratory research, clinical translation, technology transfer, industrial scale-up and international collaboration. The article does not argue that biomedical engineering is inherently WMD-related. Rather, it argues that certain capabilities may acquire indirect dual-use relevance when amplified through a series of capability multipliers, including tacit knowledge, clinical translation, regulatory science, manufacturing systems and global collaboration. By shifting analytical attention from material transfer to capability formation, the article proposes a lifecycle governance framework that embeds awareness, institutional responsibility, due diligence and responsible commercialisation across the biomedical innovation

pathway. It argues that future implementation of UNSCR 1540 should complement rather than replace existing material-centric approaches and should preserve scientific openness through a principle of responsible openness: maintaining innovation ecosystems that are as open as possible and as closed as necessary.

1 Introduction

United Nations Security Council Resolution (UNSCR) 1540, adopted in 2004, represents one of the cornerstones of the contemporary non-proliferation architecture. Unlike traditional arms control agreements, which primarily regulate the behaviour of States, UNSCR 1540 established binding obligations on all States to prevent non-State actors from acquiring, developing, manufacturing, possessing, transporting, transferring or using nuclear, chemical and biological weapons, their means of delivery and related materials. Over the past two decades, implementation efforts have concentrated largely on domestic legislation, export controls, strategic trade management, border security and the physical protection of sensitive materials. These measures have made important contributions to international security and remain indispensable components of the global non-proliferation regime.

However, rapid advances in science and technology are increasingly challenging assumptions underlying conventional non-proliferation governance. Biomedical engineering, synthetic biology, artificial intelligence, automation and digital research infrastructures are transforming the ways in which scientific capabilities are generated, organised and disseminated. At the same time, universities, research hospitals, translational centres, start-up ecosystems and international scientific networks have become central sites of knowledge production and technological innovation. Although these developments have generated extraordinary benefits for healthcare and society, they also raise broader questions concerning whether material-centric approaches alone remain sufficient to address emerging forms of dual-use risk.

These questions acquire renewed significance in light of Resolution 2663 (2022), which extended the mandate of the 1540 Committee to 2032 and explicitly recognised the importance of scientific and technological developments and evolving proliferation risks (United Nations Security Council 2022; United Nations Security Council 1540 Committee 2025). With a comprehensive review of UNSCR 1540 implementation scheduled for 2027, there is a timely opportunity to reconsider how existing implementation frameworks might adapt to increasingly complex innovation ecosystems. Yet much of the existing literature on UNSCR 1540, research security and dual-use governance continues to focus on technologies, materials and knowledge transfer, while paying comparatively less attention to the processes through which capabilities are cultivated, amplified and diffused across scientific and industrial systems.

This article argues that biomedical engineering provides an illustrative case through which to examine these challenges. Biomedical engineering is not inherently associated with weapons of mass destruction, nor does this article suggest that biomedical education or research should be viewed primarily through a security lens. Rather, biomedical engineering represents a highly interdisciplinary innovation ecosystem spanning education, doctoral training, laboratory research, clinical translation, technology transfer, industrial scale-up and international collaboration. Within such ecosystems, capabilities are not simply transferred through access to information or materials; they are progressively formed, amplified and diffused through tacit knowledge, clinical practice, regulatory science, manufacturing systems and organisational networks. Consequently, the challenge facing future UNSCR 1540 implementation is not one of restricting scientific openness, but of understanding how proportionate governance can be embedded across the lifecycle of capability formation while preserving the benefits of international scientific collaboration.

Accordingly, this article asks: *How should UNSCR 1540 address capability formation across biomedical engineering education and innovation ecosystems?* It proposes that capability formation offers a useful analytical lens for understanding upstream dimensions of non-proliferation governance. Building upon insights from the literature on tacit knowledge, technological capability and dual-use governance, the article develops a framework centred on five interconnected concepts: capability formation, capability multipliers, capability diffusion, lifecycle governance and responsible openness. It argues that capabilities are amplified rather than merely transferred, and that effective non-proliferation in the twenty-first century increasingly requires attention to the institutions and processes through which capabilities are generated, translated and diffused.

The article makes three principal contributions. First, it introduces capability formation as an analytical perspective for understanding upstream dimensions of non-proliferation governance. Second, it conceptualises tacit knowledge, clinical translation, technology transfer, quality systems, manufacturing and international collaboration as capability multipliers that amplify knowledge into practical and scalable capacities. Third, it proposes a lifecycle governance framework that embeds awareness, institutional responsibility, due diligence and responsible commercialisation across the biomedical innovation pathway. Rather than advocating greater restrictions on science, the article advances a philosophy of responsible openness, seeking to ensure that biomedical innovation ecosystems remain, in the words increasingly adopted in international science policy, "as open as possible, and as closed as necessary". To ground this framework, the article uses the medical radioisotope ecosystem, particularly the Mo-99/Tc-99m supply chain and radioactive-source lifecycle governance, as a running illustrative case. This case is selected because it is overwhelmingly beneficial to healthcare while also demonstrating how capabilities, materials, institutions, and governance mechanisms intersect across education, clinical practice, industrial production, and international supply chains.

The overall argument developed throughout this article is summarised in Figure 1. It conceptualises biomedical innovation as a continuum extending from classroom education to global mobility, through which capabilities are progressively formed, amplified and diffused. The figure also highlights the article's central proposition: that potential dual-use relevance emerges not from biomedical engineering itself, but from the accumulation and amplification of capabilities across innovation ecosystems. Consequently, future implementation of UNSCR 1540 should complement existing material-centric mechanisms with proportionate lifecycle governance while preserving the openness on which scientific progress ultimately depends.

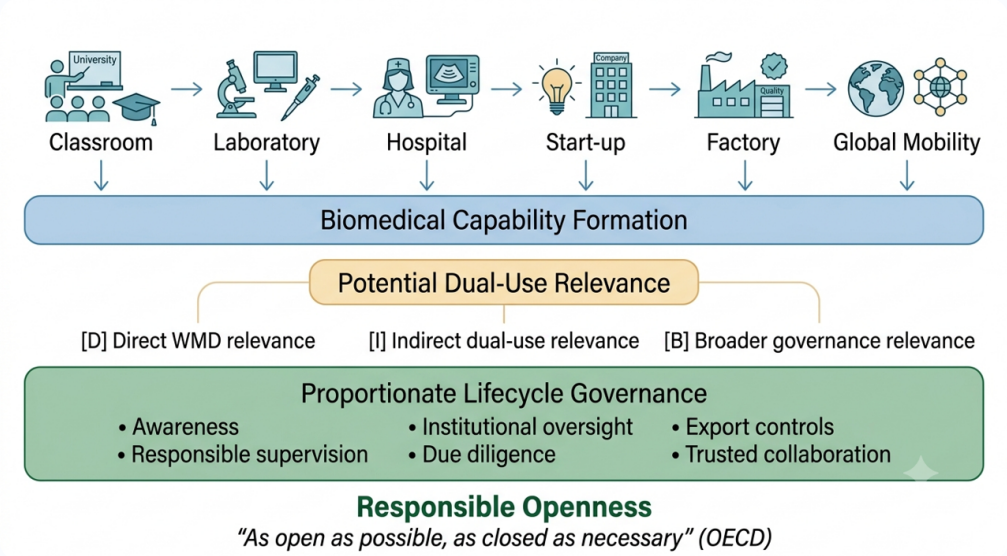

Figure 1: From classroom to commercialisation: biomedical capability formation pathways and implications for UNSCR 1540 governance. Biomedical capabilities are progressively formed across education, research, clinical translation, entrepreneurship, manufacturing and international collaboration. Potential dual-use relevance emerges primarily through the accumulation and amplification of capabilities, highlighting the need for proportionate lifecycle governance while preserving responsible openness.

To ground this framework, the article uses the medical radioisotope ecosystem as a running illustrative case, showing how beneficial biomedical capabilities are formed, amplified and governed across educational, clinical, industrial and international settings.

2 UNSCR 1540 and the Material-Centric Paradigm

2.1 Origins and Objectives

United Nations Security Council Resolution (UNSCR) 1540 was unanimously adopted in April 2004 in response to growing concerns that non-State actors might acquire nuclear, chemical or biological weapons and their means of delivery (United Nations Security Council 2004). Unlike traditional arms control treaties, which primarily regulate relations among States, UNSCR 1540 established legally binding obligations on all Member States to refrain from supporting non-State actors seeking to develop, acquire, manufacture, possess, transport, transfer or use weapons of mass destruction (WMD), and to adopt appropriate domestic measures to prevent such activities. These obligations include the enactment of national legislation and the establishment of effective controls over related materials (United Nations Security Council 2004; United Nations Office for Disarmament Affairs 2024; Crail 2006).

UNSCR 1540 emerged against a broader post-Cold War concern that terrorist groups and other non-State actors might exploit globalisation, technological diffusion and transnational supply chains to gain access to WMD-related capabilities. Consequently, the Resolution sought to strengthen international security by extending non-proliferation responsibilities beyond States and by establishing a comprehensive framework encompassing prevention, control and enforcement (Crail 2006).

Over the subsequent two decades, UNSCR 1540 has become an integral component of the international non-proliferation architecture. Together with existing treaty regimes and export control arrangements, it has contributed to strengthening national implementation mechanisms and promoting international cooperation in countering proliferation risks (United Nations Office for Disarmament Affairs 2024).

2.2 Traditional Implementation Mechanisms

In practice, implementation of UNSCR 1540 has focused predominantly on controlling physical materials, sensitive technologies and identifiable transfers. Core measures typically include export controls, strategic trade management, border security, customs enforcement, licensing systems and physical protection of sensitive facilities and materials (United Nations Security Council 2004). National implementation matrices developed by the 1540 Committee similarly emphasise legislation, accounting and control measures, border controls and law-enforcement mechanisms.

Such approaches reflect the historical origins of non-proliferation governance, which developed primarily around the control of fissile materials, chemical precursors, biological agents and delivery systems. Export control regimes, end-user verification and international assistance programmes have therefore become central tools for reducing opportunities for proliferation by limiting access to controlled goods and technologies (Crail 2006).

This material-centric orientation has proven highly effective in addressing many traditional proliferation challenges. Physical materials remain indispensable for the development of nuclear, chemical and biological weapons, and robust export control systems continue to constitute a cornerstone of international non-proliferation efforts. Consequently, recognising emerging governance challenges should not be interpreted as diminishing the continuing importance of material controls.

2.3 Material-Centric Assumptions

Underlying existing implementation practices are several implicit assumptions. First, proliferation risks are generally assumed to arise primarily from access to sensitive materials and technologies. Second, knowledge and capabilities are often treated as transferable entities that accompany physical goods or explicit technical information. Third, governance interventions are typically concentrated at points of transfer, such as exports, imports or cross-border transactions.

These assumptions have historically provided a practical basis for implementation. However, they are less well suited to understanding forms of capability that reside in people, institutions and organisational processes. Studies of tacit knowledge and technological capability suggest that the ability to reproduce complex systems depends not only on information, but also on experience, organisational routines and accumulated expertise (MacKenzie and Spinardi 1996; Ben Ouagrham-Gormley 2014; Revill 2014). Consequently, access to information alone does not necessarily translate into operational capability.

Moreover, contemporary innovation ecosystems increasingly rely upon distributed networks of universities, hospitals, start-ups, manufacturers and international collaborations. In such contexts, capabilities are progressively cultivated and amplified through training, clinical practice, regulatory experience and industrial scale-up rather than simply transferred through isolated transactions. These developments do not invalidate material-centric approaches, but they suggest that existing implementation frameworks may capture only part of the processes through which strategically relevant capabilities emerge.

2.4 Resolution 2663 and Emerging Governance Gaps

Recognition of these challenges has become increasingly visible in recent discussions surrounding the future implementation of UNSCR 1540. Resolution 2663 (2022), which extended the mandate of the 1540 Committee to 2032, explicitly acknowledged the importance of scientific and technological developments and evolving proliferation risks (United Nations Security Council 2022). The resolution also established a comprehensive review process scheduled for 2027, creating an opportunity to reassess the adequacy of existing implementation approaches.

Rapid advances in artificial intelligence, synthetic biology, automation and digital research infrastructures have further highlighted the increasingly distributed nature of scientific capability formation (Organisation for Economic Co-operation and Development 2025; World Health Organization 2022). At the same time, concerns surrounding research security, international collaboration and dual-use technologies have expanded discussions beyond traditional material controls. Nevertheless, many of these debates continue to focus on technologies and knowledge transfer, while paying comparatively less attention to the broader innovation ecosystems through which capabilities are generated, amplified and diffused (United Nations Security Council 2022; United Nations Security Council 1540 Committee 2025; Organisation for Economic Co-operation and Development 2025).

This article does not suggest that UNSCR 1540 should evolve into a general framework for governing scientific knowledge, nor does it advocate expanding the legal scope of the Resolution itself. Rather, it argues that Resolution 2663 and the forthcoming 2027 comprehensive review provide an opportunity to consider how implementation practices might better recognise capability formation across biomedical innovation ecosystems. Such an approach would complement rather than replace existing material-centric mechanisms and would seek to preserve scientific openness while addressing emerging forms of dual-use relevance.

Accordingly, the following section turns to biomedical engineering as an illustrative case of capability formation, examining how education, doctoral training, clinical translation, entrepreneurship, industrial scale-up and international collaboration collectively contribute to the generation and diffusion of scientific and technological capabilities. The key gap, therefore, is not that UNSCR 1540 ignores technology, but that its implementation architecture remains better suited to controlling materials and transfers than to recognising how capabilities are formed across innovation ecosystems.

3 From Classroom to Commercialisation: Biomedical Engineering as a Capability Formation Ecosystem

3.1 Biomedical Engineering as a Convergence Discipline

Biomedical engineering occupies a distinctive position at the intersection of engineering, life sciences and medicine. Unlike disciplines organised around a single scientific domain, biomedical engineering integrates knowledge from biology, physics, electronics, computer science, materials science, mechanics and clinical medicine to address healthcare challenges. Recent studies have increasingly characterised biomedical engineering as part of a broader convergence paradigm in which traditionally distinct disciplines interact to generate new forms of innovation (National Research Council 2014; Yock et al. 2015). Consequently, biomedical engineering should not be understood solely as a collection of technologies or products, but rather as an innovation ecosystem linking education, research, clinical practice, entrepreneurship and manufacturing.

This ecosystem perspective is particularly relevant to contemporary discussions surrounding dual-use governance. Whereas traditional non-proliferation approaches focus primarily on materials and technologies, biomedical engineering highlights the importance of institutions, professional communities and translational processes in generating capabilities. The pathways through which capabilities emerge therefore extend far beyond isolated technologies and encompass a continuum of activities stretching from education to industrial deployment.

3.2 Knowledge Formation

Capability formation begins long before technologies enter industrial production. Undergraduate education provides the foundations of technical literacy, engineering judgement and interdisciplinary problem-solving. Postgraduate education and doctoral training further deepen these capabilities through laboratory practice, experimental design, troubleshooting and supervision. Importantly, many aspects of scientific competence cannot be fully codified in publications or manuals but are acquired through apprenticeship, experience and interaction with research communities (MacKenzie and Spinardi 1996; Revill 2014).

Biomedical engineering education is particularly characterised by its interdisciplinary nature. Students and researchers routinely combine engineering methods with biological systems, clinical requirements and computational approaches. As a result, doctoral training serves not merely as a mechanism for knowledge dissemination but as a process through which tacit skills, practical experience and professional norms are accumulated. These forms of knowledge constitute important foundations upon which subsequent innovation activities are built.

At this stage, however, capabilities remain largely embedded within individuals and research groups. They become more consequential when they are translated into reproducible systems, organisational practices and scalable technologies.

3.3 Capability Amplification

The transformation of scientific knowledge into practical capability depends on processes that extend beyond the laboratory. Clinical translation plays a particularly important role in biomedical engineering because healthcare technologies must operate within complex clinical environments involving patients, clinicians, hospitals and regulatory systems. Interactions between engineers and healthcare professionals generate procedural knowledge and contextual understanding that cannot easily be derived from publications alone. Translational infrastructures such as Biomedical Research Centres and collaborative clinical networks therefore serve as important mechanisms through which laboratory discoveries are transformed into clinically relevant innovations (National Institute for Health and Care Research 2024a, 2024b).

Technology transfer and entrepreneurship constitute additional stages of capability amplification. Universities increasingly engage in patenting, licensing and spin-out formation, creating pathways through which academic discoveries are converted into commercial products and services (Association of University Technology Managers 2024; UK Research and Innovation 2024). Through these processes, scientific knowledge becomes embedded within organisations, intellectual property portfolios and entrepreneurial teams.

Industrial scale-up further amplifies capabilities. Manufacturing, quality management systems and regulatory science transform prototypes into robust, reproducible and scalable products. Standards such as ISO 13485 and evolving regulatory frameworks exemplify how industrial capabilities emerge through systematic processes rather than isolated technological breakthroughs (ISO 13485 2016; U.S. Food and Drug Administration 2024). Consequently, capabilities are progressively amplified as they move from laboratory environments into clinical, commercial and industrial settings.

3.4 Capability Diffusion

Capabilities are not confined to individual institutions. Scientific collaboration, professional mobility and international networks contribute to their diffusion across organisational and geographical boundaries. Conferences, joint projects, researcher exchanges and industrial partnerships facilitate the circulation of expertise and promote innovation. These mechanisms have been central to the remarkable successes of contemporary biomedical engineering and are indispensable to scientific progress.

However, capability diffusion also raises broader governance questions. The movement of people, organisational practices and collaborative networks can contribute to the dissemination of knowledge and expertise in ways that are difficult to capture through traditional export controls alone. Trusted research initiatives and responsible international collaboration frameworks increasingly recognise the importance of balancing openness with appropriate risk awareness (National Protective Security Authority 2024a).

Importantly, capability diffusion should not be interpreted primarily as a security problem. International collaboration remains fundamental to scientific advancement and global health. Rather, diffusion highlights the limitations of viewing proliferation solely through the lens of material transfer and suggests the need for governance approaches that recognise the social and institutional dimensions of capability generation.

3.5 Illustrative Case: The Medical Radioisotope Ecosystem

The medical radioisotope ecosystem provides a concrete illustration of the classroom-to-commercialisation pathway described in Figure 1. Education in medical physics, radiochemistry and biomedical engineering provides the codified foundations of radiation physics, dosimetry, radioactive decay and source categorisation. These foundations become operationally meaningful only through supervised laboratory practice, where researchers and technicians acquire tacit skills in radiopharmaceutical preparation, isotope handling, contamination control and quality assurance. In hospitals, these capabilities are embedded within nuclear medicine and radiotherapy workflows involving clinicians, radiopharmacists, medical physicists and radiation safety officers. Beyond the hospital, radioisotope capability depends on commercial production, quality systems, transport logistics and international supply chains. The Mo-99/Tc-99m pathway is particularly illustrative because Tc-99m is widely used in nuclear medicine while depending on a technically complex and time-sensitive international supply chain (OECD Nuclear Energy Agency 2019, 2023). The significance of this case lies not in any individual technology or institution, but in the progressive formation and diffusion of capability across educational, clinical, industrial and international settings.

3.6 Biomedical Innovation Ecosystems as Upstream Sites of Governance

Taken together, these observations suggest that biomedical engineering should be understood as a capability formation ecosystem rather than a collection of discrete technologies. Capabilities are progressively generated through knowledge formation, amplified through clinical translation and industrial development, and diffused through international networks and mobility. The significance of these processes lies not in implying that biomedical engineering is inherently associated with weapons of mass destruction, but in recognising that innovation ecosystems constitute upstream sites where capabilities emerge and evolve.

This perspective does not replace material-centric approaches to non-proliferation. Rather, it complements them by shifting analytical attention towards the processes through which capabilities are cultivated. Understanding these processes provides a foundation for examining why capabilities are amplified rather than merely transferred. The following section, therefore, develops the concept of capability multipliers and explores the mechanisms through which biomedical capabilities become increasingly practical, reproducible and scalable.

4 Capability Multipliers: Why Capabilities Are Amplified Rather than Merely Transferred

Whereas the previous section identified the stages through which biomedical capabilities are formed, this section examines the mechanisms that amplify them. The preceding section argued that biomedical engineering should be understood as a capability formation ecosystem extending from education to industrial translation. However, capability formation alone does not explain why certain forms of expertise become practical, reproducible and scalable. Existing non-proliferation frameworks frequently assume that capabilities accompany information or technologies and can therefore be controlled through material restrictions and transfer mechanisms. Yet studies of technological capability and tacit knowledge suggest that capabilities are not merely transferred but progressively amplified through organisational, institutional and social processes (MacKenzie and Spinardi 1996; Ben Ouagrham-Gormley 2014; Revill 2014).

This section develops the concept of *capability multipliers* to explain the mechanisms through which capabilities become increasingly operational and influential. As illustrated in Figure 2, biomedical capabilities are amplified through a series of interconnected processes involving tacit knowledge, clinical translation, technology transfer, quality systems, manufacturing and international collaboration. Together, these mechanisms transform technical knowledge into robust and scalable capabilities.

4.1 Tacit Knowledge as a Capability Multiplier

The importance of tacit knowledge has long been recognised in studies of technological capability and proliferation. MacKenzie and Spinardi argued that complex technologies cannot be reproduced solely from codified information because successful implementation depends upon practical experience, organisational routines and engineering judgement (MacKenzie and Spinardi 1995; Ben Ouagrham-Gormley 2014; Revill 2014). Similarly, studies of biological weapons programmes demonstrate that access to scientific publications or equipment does not automatically translate into operational capability. Instead, capabilities depend upon tacit expertise accumulated through training, experimentation and organisational learning (Ben Ouagrham-Gormley 2014).

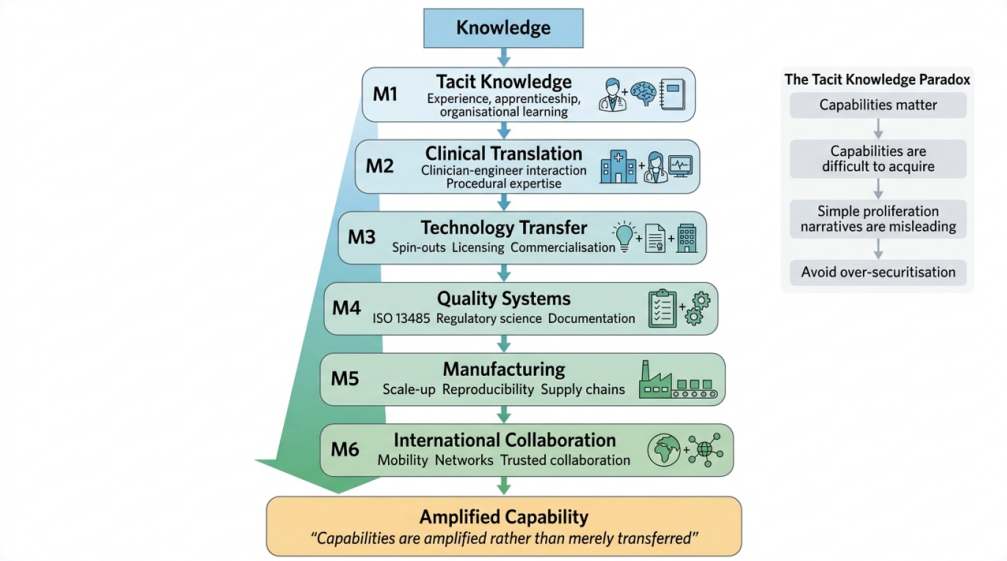

Figure 2: Capability multipliers across the biomedical innovation ecosystem. Biomedical capabilities are amplified rather than merely transferred through a series of interconnected multipliers, including tacit knowledge, clinical translation, technology transfer, quality systems, manufacturing and international collaboration. These mechanisms progressively transform knowledge into practical, reproducible and scalable capabilities. The tacit knowledge paradox highlights that while capabilities matter, they are difficult to acquire, cautioning against simplistic proliferation narratives and supporting proportionate approaches to governance.

Within biomedical engineering, tacit knowledge is acquired through laboratory practice, troubleshooting, iterative experimentation and supervision. Such capabilities are difficult to codify and are often embedded within research groups and professional communities. Consequently, the acquisition of knowledge does not necessarily imply the acquisition of capability. Rather, tacit knowledge acts as a multiplier that converts information into practical competence.

4.2 Clinical Translation as a Capability Multiplier

Clinical translation represents one of the most distinctive features of biomedical engineering and constitutes a capability multiplier largely absent from traditional discussions of dual-use governance. Translational processes require interactions between engineers, clinicians, patients and healthcare systems. Through these interactions, technologies are refined, validated and adapted to clinical workflows, generating procedural knowledge and operational experience that extend beyond laboratory research.

Institutions such as Biomedical Research Centres and translational research collaborations provide environments in which scientific discoveries are transformed into clinically useful innovations (National Institute for Health and Care Research 2024a, 2024b). Clinical translation therefore amplifies capabilities by integrating technical knowledge with human factors, regulatory requirements and practical experience. In this sense, hospitals and translational infrastructures function not merely as end-users of technology, but as sites where capabilities are enhanced and institutionalised.

4.3 Technology Transfer and Regulatory Science

Technology transfer constitutes another mechanism through which capabilities are amplified. Universities increasingly rely on technology transfer offices, licensing activities and spin-out companies to convert research outputs into products and services (Association of University Technology Managers 2024; UK Research and Innovation 2024). Through these processes, scientific knowledge becomes embedded within organisational structures, entrepreneurial teams and intellectual property frameworks.

Regulatory science further contributes to capability amplification. Product development requires validation, documentation, risk management and compliance with regulatory standards. These activities generate forms of organisational knowledge that enable technologies to move from proof-of-concept demonstrations to practical implementation. Consequently, technology transfer and regulatory science transform isolated discoveries into sustainable innovation platforms.

4.4 Quality Systems and Manufacturing

Industrial scale-up provides another important multiplier. Quality management systems, manufacturing processes and regulatory standards enable reproducibility, reliability and scalability. Standards such as ISO 13485 and the United States Food and Drug Administration's Quality Management System Regulation provide structured mechanisms through which design, production and quality assurance are integrated into industrial practice (ISO 13485 2016; U.S. Food and Drug Administration 2024).

Manufacturing capability extends beyond physical production. It encompasses process control, documentation, supply chains and organisational routines that permit technologies to be deployed at scale. Consequently, manufacturing transforms prototypes into mature systems and represents a key mechanism through which capabilities become increasingly robust and transferable.

4.5 International Collaboration and Mobility

Scientific collaboration and international mobility have historically played central roles in biomedical innovation. Conferences, joint research programmes, mobility schemes and industrial partnerships facilitate the circulation of expertise and promote scientific progress. International networks therefore act as multipliers by accelerating learning and enabling access to complementary knowledge and resources.

Trusted research frameworks increasingly recognise that collaboration itself can contribute to capability diffusion and should therefore be accompanied by proportionate risk awareness and institutional responsibility (National Protective Security Authority 2024a). Nevertheless, collaboration remains fundamental to scientific advancement and should not be interpreted primarily through a security lens. Rather, international networks should be understood as mechanisms that amplify capabilities while simultaneously requiring appropriate governance.

4.6 The Tacit Knowledge Paradox

The medical radioisotope ecosystem also illustrates how capability multipliers operate in practice. Tacit knowledge is required to handle sealed and unsealed radioactive sources, operate technetium generators, conduct radiolabelling under time pressure and respond to contamination events. Clinical translation amplifies these skills by embedding them within nuclear medicine workflows, where technical competence is coordinated with patient care, quality assurance and radiation protection. Technology transfer and commercialisation can then convert radiopharmaceutical formulations, generator systems or quality-control instrumentation into organisational assets. Quality systems, including Good Manufacturing Practice and medical-device quality management, further transform laboratory-scale practice into reproducible and auditable procedures. Manufacturing and international collaboration add additional layers of capability concentration and diffusion, as Mo-99/Tc-99m supply depends on specialised production facilities, transport networks and international coordination (OECD Nuclear Energy Agency 2019, 2023). This example reinforces the article's central claim: capabilities are amplified through institutional settings rather than merely transferred as information.

The preceding discussion reveals a paradox at the heart of capability formation. On the one hand, capabilities matter because practical competence depends upon tacit knowledge, organisational learning and institutional experience. This suggests that capabilities deserve greater attention within contemporary non-proliferation discussions. On the other hand, the importance of tacit knowledge also implies that capabilities are difficult to acquire, reproduce and weaponise. Access to information alone is insufficient to create operational capabilities, and simple narratives equating scientific knowledge with proliferation risk are therefore misleading.

This *tacit knowledge paradox* has important implications for governance. It suggests that the recognition of capability formation should not lead to excessive securitisation of science. Instead, it reinforces the need for proportionate approaches that acknowledge both the significance and the complexity of capability acquisition. In this sense, the existence of capability multipliers strengthens rather than weakens the case for responsible openness.

Taken together, these observations suggest that capabilities are amplified rather than merely transferred. Biomedical capabilities emerge through a network of social, organisational and institutional mechanisms that progressively transform knowledge into practical and scalable capacities. Understanding these multipliers provides a bridge between capability formation and governance, because different stages of amplification create distinct opportunities for oversight and responsibility. The next section therefore examines how these capabilities relate to weapons of mass destruction and explores the varying forms of direct, indirect and broader relevance associated with biomedical innovation ecosystems.

5 Dual-Use Relevance and Potential WMD Pathways

The preceding sections have argued that biomedical engineering capabilities are progressively formed and amplified through interconnected innovation ecosystems. However, recognising capability formation does not imply that biomedical engineering is inherently associated with WMD. Rather, the relationship between biomedical capabilities and proliferation risks is heterogeneous and context-dependent. Distinguishing different forms of relevance is therefore essential to avoid both underestimating emerging challenges and over-securitising legitimate scientific activity.

Table 1 provides a conceptual framework distinguishing direct WMD relevance [D], indirect dual-use relevance [I], and broader governance relevance [B]. Direct relevance refers to activities or materials with explicit connections to WMD acquisition or development. Indirect relevance captures enabling technologies and capabilities that may acquire dual-use significance under specific circumstances. Broader governance relevance refers to institutions, ecosystems and processes that are important for understanding capability formation but do not themselves constitute proliferation risks.

Table 1: Direct, indirect and broader relevance to UNSCR 1540 across biomedical capability formation. [D] = direct WMD relevance; [I] = indirect dual-use relevance; [B] = broader governance relevance. ❍ = low relevance; ● = moderate relevance; ●● = high relevance. Most biomedical engineering activities possess indirect dual-use or broader governance relevance rather than direct WMD relevance.

Domain	[D]	[I]	[B]
Biomedical engineering education	❍	●	●●
Doctoral training and tacit knowledge	❍	●	●●
Clinical translation and hospitals	❍	●	●●
Technology transfer and start-ups	❍	●	●●
Medical radioactive sources	●●	●	●●
Synthetic biology and biofoundries	❍	●●	●●
AI, robotics and automation	❍	●●	●●
Quality systems and manufacturing	❍	●	●●
International mobility and collaboration	❍	●	●●

5.1 Biological Pathways

The life sciences have long been recognised as possessing dual-use characteristics. Advances in synthetic biology, bioinformatics, laboratory automation and biological engineering have expanded the range of capabilities available to researchers and industries (National Academies of Sciences, Engineering, and Medicine 2018; World Health Organization 2022; Jefferson et al. 2014). However, the overwhelming majority of these activities are directed towards beneficial purposes, including diagnostics, therapeutics, vaccine development and public health.

From the perspective of UNSCR 1540, biological pathways are relevant primarily because some enabling capabilities could, under specific circumstances, contribute to the acquisition or production of biological agents. Such concerns have motivated international efforts to strengthen responsible conduct and biosafety governance while preserving the benefits of scientific openness (World Health Organization 2022; Fink 2004). Importantly, biological capability should not be equated with biological weapons capability. Studies of tacit knowledge and organisational complexity demonstrate that the successful development of biological weapons requires far more than access to information or equipment (Ben Ouagrham-Gormley 2014; Revill 2014).

5.2 Chemical Pathways

Biomedical engineering also intersects with chemical sciences through pharmaceuticals, biomaterials, toxicology and chemical engineering. Although certain forms of chemical expertise may possess indirect dual-use significance, most biomedical applications remain far removed from chemical weapons programmes. In this context, chemical pathways should be understood primarily through the lens of enabling capabilities rather than direct weaponisation.

Consequently, chemical relevance under UNSCR 1540 is generally indirect and conditional. Regulatory oversight, responsible innovation and established chemical safety frameworks remain important mechanisms for managing such risks without imposing unnecessary restrictions on legitimate scientific activities.

5.3 Radiological Pathways

Radiological pathways provide the clearest case in which biomedical innovation ecosystems may intersect directly with UNSCR 1540 concerns. Hospitals, research institutions and medical industries routinely use radioactive sources for diagnosis, radiotherapy, sterilisation and research. These applications are overwhelmingly beneficial, but high-activity sealed sources may pose safety and security risks if they are abandoned, lost, stolen or inadequately controlled. The IAEA categorisation system therefore provides a risk-based framework for graded regulatory control of radioactive sources (International Atomic Energy Agency 2005a).

The 1987 Goiânia accident illustrates this point without implying that nuclear medicine itself is a proliferation threat. A Cs-137 teletherapy source from a former radiotherapy facility was left outside effective institutional control, removed from its shielding and dispersed, causing severe radiological consequences (International Atomic Energy Agency 1988). The governance failure was not located in medical physics education, clinical practice or radiochemistry as such, but in the breakdown of institutional accountability for a high-activity source after it ceased to be actively managed by a licensed institution. In this sense, direct relevance emerges not from knowledge itself, but from the point at which a physical high-activity source becomes separated from institutional accountability.

This distinction is central to the argument of this article. The educational and clinical capabilities underpinning nuclear medicine are primarily of indirect dual-use or broader governance relevance. Direct material relevance is concentrated at specific lifecycle nodes, including high-activity sealed sources, transport, storage and end-of-life management. Consequently, the appropriate governance response is not to restrict biomedical capability formation, but to strengthen source accountability, institutional continuity and lifecycle management (International Atomic Energy Agency 2003, 2018, 2019).

5.4 Means of Delivery and Enabling Capabilities

Beyond biological, chemical and radiological dimensions, biomedical engineering contributes to a range of enabling technologies, including robotics, automation, sensing, microfluidics and artificial intelligence. These technologies are not weapons in themselves and are rarely developed with malicious intent. However, they may contribute indirectly to delivery systems, process optimisation or automation capabilities under certain circumstances.

Such examples are better understood as instances of indirect dual-use relevance rather than direct proliferation pathways. This distinction is particularly important because the value of these technologies for medicine and public health overwhelmingly outweighs hypothetical misuse scenarios. Consequently, governance approaches should focus on proportionality and context rather than broad restrictions.

5.5 Avoiding Over-Securitisation

Recognition of dual-use relevance should not lead to the securitisation of biomedical engineering as a whole. Dual-use research is not synonymous with dual-use research of concern (DURC), and most biomedical education, research and innovation activities possess little or no direct relevance to WMD acquisition (Uhlenhaut et al. 2012; Bruggen 2012; McLeish and Nightingale 2007). Overly expansive interpretations of dual-use risk risk undermining academic freedom, scientific collaboration and the societal benefits generated by open research.

Moreover, the importance of tacit knowledge itself suggests that simple narratives linking access to information with proliferation are inadequate. Capabilities are difficult to acquire, organisationally complex and embedded within institutions and communities of practice. This tacit knowledge paradox reinforces the need for proportionate approaches to governance and argues against treating scientific openness primarily as a security liability.

Accordingly, the significance of biomedical engineering to UNSCR 1540 lies less in direct WMD relevance than in broader questions concerning how capabilities are formed, amplified and diffused. Recognising these distinctions provides an essential foundation for the next section, which develops a lifecycle governance framework aimed at embedding awareness, responsibility and due diligence across biomedical innovation ecosystems while preserving responsible openness.

6 Towards a Biomedical Capability Formation Lifecycle Governance Framework

The preceding sections have argued that biomedical capabilities are progressively formed, amplified and diffused through interconnected innovation ecosystems. They have also demonstrated that most biomedical engineering activities possess indirect dual-use or broader governance relevance rather than direct relevance to WMD. Consequently, the challenge facing future implementation of UNSCR 1540 is not one of expanding material controls indiscriminately, but of identifying appropriate points at which governance may be embedded throughout the capability lifecycle.

Building upon this perspective, Table 2 summarises the principal stages of biomedical capability formation, the actors involved, the capabilities generated and corresponding governance responses. Rather than concentrating interventions solely at points of transfer, this framework embeds governance across education, research, clinical translation, entrepreneurship, manufacturing and international collaboration.

Table 2: Biomedical Capability Formation Matrix. Different stages of the biomedical innovation ecosystem contribute to capability formation and diffusion and are associated with distinct governance responses. TTO = technology transfer office; dual-use relevance is indicated qualitatively and refers primarily to indirect or broader relevance rather than direct WMD relevance.

Stage	Typical actors	Capability formed	Relevance	Governance response
Classroom	Students, educators	Technical literacy and engineering judgement	Low	Ethics and dual-use awareness
Laboratory	PhD researchers, supervisors	Tacit knowledge and experimental practice	Moderate	Responsible supervision and research culture
Hospital	Clinician-engineer teams	Procedural expertise and clinical translation	Moderate	Clinical governance and oversight
Start-up	Founders, TTOs, investors	Productisation and commercialisation	Moderate	Technology-transfer due diligence
Factory	Manufacturers and industry	Manufacturing, QMS and scale-up	High	Export compliance and end-user screening
Global mobility	International networks	Capability diffusion	High	Trusted collaboration and capacity-building

Figure 3 provides a graphical representation of this lifecycle governance framework. The objective is not to securitise biomedical innovation, but to integrate awareness, responsibility and due diligence throughout the innovation ecosystem while preserving openness and international collaboration.

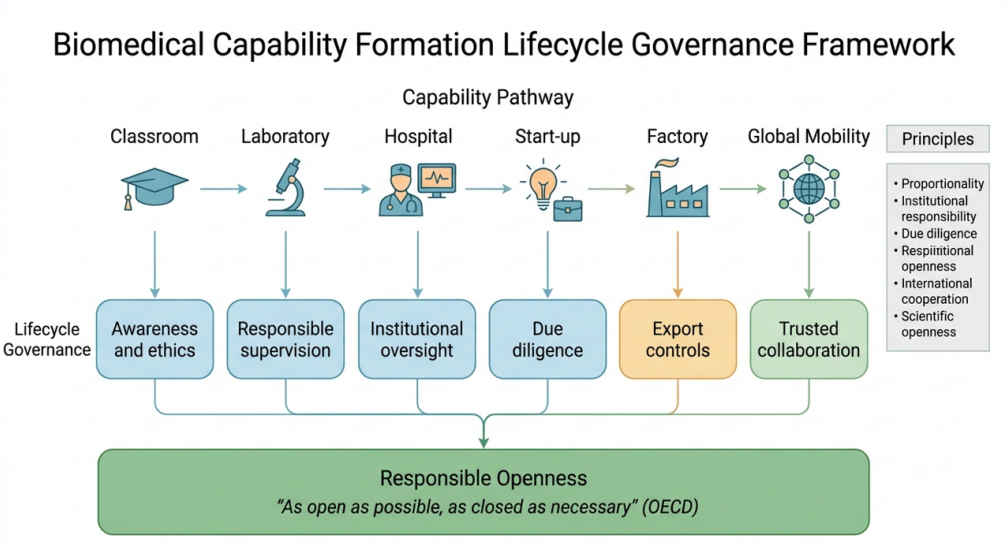

Figure 3: Biomedical capability formation lifecycle governance framework. Governance interventions are embedded throughout the biomedical innovation pathway, from education and research to commercialisation and international collaboration. Rather than concentrating solely on material controls, lifecycle governance integrates awareness, responsible supervision, institutional oversight, due diligence, export controls and trusted

collaboration in a manner consistent with the principle of responsible openness. Adapted to biomedical innovation ecosystems, this framework seeks to preserve scientific openness while addressing evolving forms of dual-use relevance under UNSCR 1540.

6.1 Awareness and Education

Capability formation begins in educational environments. Universities and professional programmes constitute the earliest stages at which awareness of dual-use issues, research integrity and responsible innovation can be introduced. Such measures should not be interpreted as attempts to transform biomedical curricula into security programmes. Rather, they provide opportunities to cultivate ethical reflection and awareness of broader societal responsibilities. Educational interventions therefore represent proportionate and low-cost mechanisms for strengthening responsible innovation without undermining scientific openness.

6.2 Responsible Supervision and Research Culture

Doctoral training and laboratory practice represent important sites for the development of tacit knowledge and experimental capabilities. Supervisors and research groups play central roles in shaping research culture, mentoring practices and professional norms. Consequently, responsible supervision and institutional support structures constitute important components of lifecycle governance. These mechanisms emphasise awareness and accountability rather than restrictive oversight, thereby maintaining the collaborative character of scientific research while encouraging responsible conduct.

6.3 Institutional Oversight and Clinical Translation

Clinical translation introduces additional layers of organisational complexity involving hospitals, ethics committees, regulatory systems and multidisciplinary teams. Translational infrastructures provide opportunities to incorporate institutional oversight and risk assessment without imposing unnecessary barriers to innovation. Clinical governance mechanisms, already established within healthcare systems, may therefore serve as natural platforms for integrating responsible innovation practices. In this sense, hospitals and translational centres are not merely users of technology but important actors in capability governance.

6.4 Due Diligence and Responsible Commercialisation

Technology transfer, entrepreneurship and start-up formation constitute critical mechanisms through which scientific discoveries become products and services. Technology transfer offices, incubators, investors and industrial partners increasingly influence the direction and scale of innovation. Consequently, due diligence and responsible commercialisation practices represent important elements of lifecycle governance. Such measures may include partner screening, awareness of international collaborations and responsible approaches to intellectual property and licensing, while preserving incentives for innovation and entrepreneurship.

6.5 Export Controls and Manufacturing Governance

Manufacturing and industrial scale-up remain closely connected to traditional non-proliferation mechanisms. Quality systems, supply chains and production capabilities transform prototypes into reproducible and scalable technologies. Existing export control frameworks and end-user screening mechanisms therefore continue to play essential roles in preventing proliferation. Rather than replacing material-centric approaches, lifecycle governance complements them by recognising that industrial capability itself represents an important multiplier of technological capacity.

6.6 Existing Precedents: Radioactive Source Lifecycle Governance

The governance of radioactive sources provides an existing precedent for the lifecycle approach proposed in this article. The IAEA Code of Conduct on the Safety and Security of Radioactive Sources and its supplementary guidance on import and export and on disused radioactive sources demonstrate how governance can travel with materials and capabilities across institutional boundaries (International Atomic Energy Agency 2004, 2005b, 2018). Rather than focusing only on initial acquisition or cross-border transfer, these instruments emphasise source categorisation, regulatory control, inventory management, import-export coordination and end-of-life management.

National mechanisms similarly illustrate lifecycle governance in practice. The U.S. National Source Tracking System, for example, tracks Category 1 and Category 2 radioactive sources, reflecting the IAEA Code of Conduct's concern with high-risk sources (U.S. Nuclear Regulatory Commission 2024). Programmes such as the U.S. International Radiological Threat Reduction Program further show how source security can be strengthened through cooperative capacity-building, physical protection upgrades and prioritisation of high-risk sites (U.S. Government Accountability Office 2007). These mechanisms are not perfect analogues for all biomedical technologies, but they provide a concrete example of how proportionate governance can be embedded across education, institutional oversight, industrial practice, transport and end-of-life management without prohibiting beneficial medical use.

6.7 Trusted Collaboration and International Capacity-Building

International mobility and scientific collaboration are indispensable to biomedical progress. Conferences, joint projects and professional networks facilitate the exchange of knowledge and accelerate innovation. Governance should therefore seek to support trusted collaboration rather than restrict scientific openness unnecessarily. Trusted research initiatives, international norms and institutional awareness mechanisms provide examples of approaches that balance openness with appropriate risk management. International capacity-building and cooperation remain central to the spirit of UNSCR 1540 and should continue to be promoted.

6.8 Towards Responsible Openness

Taken together, these observations suggest that future implementation of UNSCR 1540 should complement traditional material-centric approaches with lifecycle governance embedded across biomedical innovation ecosystems. Such an approach does not require expanding the legal scope of the Resolution, nor does it imply that biomedical engineering should be viewed primarily through a security lens. Rather, it recognises that capabilities are progressively formed and amplified through institutions and social processes, and that governance should therefore be proportionate, context-dependent and compatible with scientific openness.

The ultimate objective is responsible openness: maintaining biomedical innovation ecosystems that are, in the words increasingly adopted in international science policy, “as open as possible, as closed as necessary” (Organisation for Economic Co-operation and Development 2025; World Health Organization 2022; National Protective Security Authority 2024b). By embedding awareness, institutional responsibility, due diligence and trusted collaboration throughout the capability lifecycle, future implementation of UNSCR 1540 can address emerging challenges while preserving the openness upon which scientific progress ultimately depends.

7 Policy Implications for the 2027 Comprehensive Review

The preceding discussion suggests that contemporary biomedical innovation ecosystems challenge the implicit assumptions underlying predominantly material-centric approaches to non-proliferation. While UNSCR 1540 remains fundamentally concerned with preventing non-State actors from acquiring weapons of mass destruction and related materials, emerging technologies and increasingly interconnected innovation systems highlight the importance of capability formation and diffusion as complementary dimensions of governance. These developments provide opportunities for reflection in the context of Resolution 2663 and the forthcoming 2027 Comprehensive Review (United Nations Security Council 2022; United Nations Security Council 1540 Committee 2025).

Importantly, recognising capability formation does not imply the need to expand the legal scope of UNSCR 1540 or to reinterpret ordinary biomedical research as inherently security-sensitive. Rather, the findings of this article suggest that existing implementation practices may benefit from greater attention to the social and institutional processes through which capabilities are cultivated and amplified.

7.1 From Material Control to Capability Awareness

Traditional implementation mechanisms under UNSCR 1540 have focused primarily on export controls, accounting systems and the physical protection of materials. While these measures remain indispensable, capability-oriented approaches emphasise the importance of awareness, institutional responsibility and responsible innovation. Universities, hospitals, technology transfer offices and industrial actors represent important sites at which awareness and proportionate governance practices may be embedded without imposing additional legal obligations.

Consequently, future implementation efforts could complement material controls with educational initiatives, institutional guidance and voluntary awareness mechanisms that encourage responsible conduct across biomedical innovation ecosystems.

A capability-based approach also shifts attention towards the meso-level institutions where governance is actually operationalised. Universities, principal investigators, hospital radiation safety committees, technology transfer offices, professional bodies, national regulators and international organisations all shape how capabilities are cultivated, supervised, commercialised and diffused. These actors sit between high-level international obligations and individual researchers. Strengthening UNSCR 1540 implementation, therefore, requires not only national legislation and border controls, but also institutional guidance, professional norms and practical due diligence mechanisms embedded within the organisations that manage biomedical innovation.

7.2 Strengthening Capacity-Building and Responsible Innovation

Resolution 2663 places renewed emphasis on assistance and capacity-building. This orientation aligns naturally with a capability-based perspective. Rather than treating scientific development and non-proliferation as competing objectives, capacity-building can support both innovation and responsible governance. Training programmes, research integrity initiatives and trusted collaboration frameworks provide examples of measures that strengthen institutional capabilities while preserving openness and international cooperation.

Particular attention may be devoted to translational infrastructures, technology transfer activities and industrial quality systems, which have traditionally received less attention within non-proliferation discussions despite their roles as capability multipliers.

7.3 Supporting Trusted International Collaboration

International scientific cooperation remains indispensable to biomedical progress and global health. Excessively restrictive approaches risk undermining the openness upon which innovation depends. Accordingly, trusted collaboration should be viewed as a mechanism for strengthening responsible research ecosystems rather than as a justification for limiting scientific exchange.

Existing trusted research initiatives and international partnerships provide examples of approaches that seek to balance openness with proportionate risk management. Such practices are particularly relevant to biomedical engineering, where multidisciplinary and international collaboration are essential components of capability formation.

7.4 Avoiding Over-Securitisation

A capability-based perspective also highlights the importance of avoiding over-securitisation. The significance of tacit knowledge and organisational complexity demonstrates that access to information alone does not automatically translate into operational capability. Consequently, broad attempts to govern scientific knowledge through security frameworks may prove both ineffective and detrimental to innovation.

Future implementation of UNSCR 1540 should therefore distinguish between direct WMD relevance, indirect dual-use relevance and broader governance relevance. Such distinctions help preserve academic freedom and international cooperation while ensuring that proportionate measures are directed towards genuinely significant risks.

7.5 Towards Responsible Openness

The 2027 Comprehensive Review offers an opportunity to reaffirm that scientific openness and non-proliferation need not be viewed as opposing objectives. Rather than shifting from material control to knowledge control, future implementation may be understood as complementing existing approaches with greater awareness of capability formation, capability amplification and lifecycle governance.

Accordingly, this article argues that biomedical innovation ecosystems should be recognised as upstream sites of capability formation and that proportionate interventions should be embedded throughout the innovation lifecycle. Such an approach supports the broader objective of responsible openness, preserving the conditions necessary for scientific progress while addressing evolving forms of dual-use relevance in the emerging technology era.

8 Conclusion

This article has argued that contemporary biomedical innovation ecosystems present challenges that are not fully captured by traditional material-centric approaches to non-proliferation. While UNSCR 1540 was originally conceived to prevent non-State actors from acquiring weapons of mass destruction, related materials and delivery systems, the increasing complexity of biomedical engineering ecosystems highlights the importance of capability formation, amplification and diffusion as complementary dimensions of governance (Crail 2006; World Health Organization 2022; Organisation for Economic Co-operation and Development 2025).

The first contribution of this article is conceptual. It proposes that biomedical engineering should be understood not as a collection of isolated technologies but as a capability formation ecosystem extending from education and laboratory research to clinical translation, entrepreneurship, manufacturing and international collaboration. This perspective shifts attention from individual technologies towards the institutional and social processes through which capabilities are generated and reproduced.

The second contribution is theoretical. Building upon insights from studies of tacit knowledge and technological capability, the article develops the concept of capability multipliers to explain why capabilities are amplified rather than merely transferred. In particular, clinical translation, technology transfer, quality systems, manufacturing and international collaboration are identified as mechanisms that progressively transform scientific knowledge into practical, reproducible and scalable capabilities. At the same time, the tacit knowledge paradox highlights that capabilities are difficult to acquire and cautions against simplistic narratives equating access to information with proliferation risk.

The third contribution is policy-oriented. Rather than advocating the expansion of the legal scope of UNSCR 1540, the article proposes a lifecycle governance framework that embeds awareness, responsible supervision, institutional oversight, due diligence, export controls and trusted collaboration across biomedical innovation ecosystems. Such an approach complements existing material-centric mechanisms while preserving the openness upon which scientific progress depends.

Importantly, the significance of biomedical engineering to UNSCR 1540 should not be overstated. Most biomedical education, research and innovation activities possess indirect dual-use or broader governance relevance rather than direct relevance to weapons of mass destruction. Recognising capability formation therefore should not lead to the securitisation of science. Instead, it reinforces the need for proportionate and context-dependent approaches that distinguish between direct WMD relevance, indirect dual-use relevance and broader governance concerns.

Looking ahead, the 2027 Comprehensive Review provides an opportunity to reconsider how implementation of UNSCR 1540 may adapt to increasingly interconnected innovation ecosystems. Future non-proliferation governance need not shift from material control to knowledge control. Rather, it may be understood as complementing existing approaches with greater awareness of capability formation and lifecycle governance.

The medical radioisotope ecosystem illustrates this argument in concentrated form. Its relevance to UNSCR 1540 lies not in any individual biomedical technology, but in the progressive formation, amplification and diffusion of capabilities across educational, clinical, industrial and international networks. At the same time, the IAEA radioactive-source security architecture demonstrates that lifecycle governance can be designed to preserve beneficial medical use while strengthening institutional accountability.

Ultimately, the challenge is not to choose between security and openness, but to sustain both. Biomedical innovation ecosystems are among the most important drivers of scientific progress and global health. Preserving their openness while addressing evolving forms of dual-use relevance requires approaches that are proportionate, internationally cooperative and institutionally grounded. In this sense, the future of UNSCR 1540 may be guided not by the pursuit of restrictive knowledge governance, but by the principle of responsible openness—maintaining science that is, in the words increasingly adopted in international science policy, “as open as possible, as closed as necessary”.

Association of University Technology Managers. 2024. *AUTM u.s. Licensing Activity Survey*. AUTM.

Ben Ouagrham-Gormley, Sonia. 2014. *Barriers to Bioweapons: The Challenges of Expertise and Organization for Weapons Development*. Cornell University Press.

Bruggen, Koos van der. 2012. “Possibilities, Intentions and Threats: Dual Use in the Life Sciences Reconsidered.” *Science and Engineering Ethics* 18 (4): 741–56. https://doi.org/10.1007/s11948-011-9266-2.

Crail, Peter. 2006. “Implementing UN Security Council Resolution 1540: A Risk-Based Approach.” *The Nonproliferation Review* 13 (2): 355–99. https://doi.org/10.1080/10736700601012193.

Fink, Gerald R., ed. 2004. *Biotechnology Research in an Age of Terrorism*. National Academies Press. https://doi.org/10.17226/10827.

International Atomic Energy Agency. 1988. *The Radiological Accident in Goiânia*. STI/PUB/815. International Atomic Energy Agency. https://www.iaea.org/publications/3684/the-radiological-accident-in-goiania.

International Atomic Energy Agency. 2003. *Security of Radioactive Sources*. IAEA-TECDOC-1355. International Atomic Energy Agency. https://www-pub.iaea.org/MTCD/Publications/PDF/te_1355_web.pdf.

International Atomic Energy Agency. 2004. *Code of Conduct on the Safety and Security of Radioactive Sources*. International Atomic Energy Agency. https://www.iaea.org/publications/6956/code-of-conduct-on-the-safety-and-security-of-radioactive-sources.

International Atomic Energy Agency. 2005a. *Categorization of Radioactive Sources*. RS-G-1.9. IAEA Safety Standards Series. International Atomic Energy Agency. https://www.iaea.org/publications/7237/categorization-of-radioactive-sources.

International Atomic Energy Agency. 2005b. *Guidance on the Import and Export of Radioactive Sources*. International Atomic Energy Agency. https://www.iaea.org/publications/7227/code-of-conduct-on-the-safety-and-security-of-radioactive-sources-guidance-on-the-import-and-export-of-radioactive-sources.

International Atomic Energy Agency. 2018. *Guidance on the Management of Disused Radioactive Sources*. International Atomic Energy Agency. https://www.iaea.org/publications/13380/guidance-on-the-management-of-disused-radioactive-sources.

International Atomic Energy Agency. 2019. *Security of Radioactive Material in Use and Storage and of Associated Facilities*. No. 11-G (Rev. 1). IAEA Nuclear Security Series. International Atomic Energy Agency. https://www.iaea.org/publications/12360/security-of-radioactive-material-in-use-and-storage-and-of-associated-facilities.

ISO 13485:2016 Medical Devices—Quality Management Systems—Requirements for Regulatory Purposes (2016).

Jefferson, Catherine, Filippa Lentzos, and Claire Marris. 2014. "Synthetic Biology and Biosecurity: Challenging the Myths." *Frontiers in Public Health* 2: 115. https://doi.org/10.3389/fpubh.2014.00115.

MacKenzie, Donald, and Graham Spinardi. 1995. "Tacit Knowledge, Weapons Design, and the Uninvention of Nuclear Weapons." *American Journal of Sociology* 101 (1): 44–99. https://doi.org/10.1086/230699.

MacKenzie, Donald, and Graham Spinardi. 1996. *Knowing Machines: Essays on Technical Change*. MIT Press.

McLeish, Caitríona, and Paul Nightingale. 2007. "Biosecurity, Bioterrorism and the Governance of Science: The Increasing Convergence of Science and Security Policy." *Research Policy* 36 (10): 1635–54. https://doi.org/10.1016/j.respol.2007.10.003.

National Academies of Sciences, Engineering, and Medicine. 2018. *Biodefense in the Age of Synthetic Biology*. National Academies Press. https://doi.org/10.17226/24890.

National Institute for Health and Care Research. 2024a. *Biomedical Research Centres*. https://www.nihr.ac.uk.

National Institute for Health and Care Research. 2024b. *Translational Research Collaborations*. https://www.nihr.ac.uk.

National Protective Security Authority. 2024a. *Trusted Research Guidance for Academia*. NPSA.

National Protective Security Authority. 2024b. *Trusted Research Guidance for Academia*. National Protective Security Authority. https://www.npsa.gov.uk/specialised-guidance/trusted-research/trusted-research-academia.

National Research Council. 2014. *Convergence: Facilitating Transdisciplinary Integration of Life Sciences, Physical Sciences, Engineering, and Beyond*. National Academies Press. https://doi.org/10.17226/18722.

OECD Nuclear Energy Agency. 2019. *The Supply of Medical Isotopes: An Economic Diagnosis and Possible Solutions*. OECD Publishing. https://doi.org/10.1787/9b326195-en.

OECD Nuclear Energy Agency. 2023. *The Security of Supply of Medical Radioisotopes: Demand and Capacity Projections for 99Mo/99mTc for the 2023–2027 Period*. NEA/NDC/R(2023)1. OECD Nuclear Energy Agency. https://www.oecd-nea.org/jcms/pl_87477/the-security-of-supply-of-medical-radioisotopes-demand-and-capacity-projections-for-99mo/99mtc-for-the-2023-2027-period.

Organisation for Economic Co-operation and Development. 2025. *OECD Science, Technology and Innovation Outlook 2025*. OECD Publishing. https://doi.org/10.1787/5fe57b90-en.

Revill, James. 2014. "Tacit Knowledge and the Biological Weapons Regime." *Science and Public Policy* 41 (5): 597–610. https://doi.org/10.1093/scipol/sct090.

Uhlenhaut, Christian, Reinhard Burger, and Lars Schaade. 2012. "Protecting Society: Biological Security and Dual-Use Dilemma in the Life Sciences." *EMBO Reports* 14 (1): 25–30. https://doi.org/10.1038/embor.2012.195.

UK Research and Innovation. 2024. *Commercialisation and Spin-Out Activities*. https://www.ukri.org.

United Nations Office for Disarmament Affairs. 2024. *United Nations Security Council Resolution 1540*. https://disarmament.unoda.org/wmd/sc1540/.

United Nations Security Council. 2004. *Resolution 1540 (2004)*. https://undocs.org/S/RES/1540(2004).

United Nations Security Council. 2022. *Resolution 2663 (2022)*. https://undocs.org/S/RES/2663(2022).

United Nations Security Council 1540 Committee. 2025. *2027 Comprehensive Review*. https://www.un.org/en/sc/1540/comprehensive-and-annual-reviews/2027-comprehensive-review.shtml.

U.S. Food and Drug Administration. 2024. *Quality Management System Regulation: Final Rule Amending the Quality System Regulation*. https://www.fda.gov/medical-devices/quality-and-compliance-medical-devices/quality-system-qs-regulationmedical-device-good-manufacturing-practices.

U.S. Government Accountability Office. 2007. *Nuclear Nonproliferation: DOE's International Radiological Threat Reduction Program Needs to Focus Future Efforts on Securing the Highest Priority Radiological Sources*. GAO-07-282. U.S. Government Accountability Office. https://www.gao.gov/products/gao-07-282.

U.S. Nuclear Regulatory Commission. 2024. *National Source Tracking System: Background*. https://www.nrc.gov/security/byproduct/ismp/nsts/background.

World Health Organization. 2022. *Global Guidance Framework for the Responsible Use of the Life Sciences: Mitigating Biorisks and Governing Dual-Use Research*. World Health Organization. https://www.who.int/publications/i/item/9789240056107.

Yock, Paul G., Stefanos Zenios, Josh Makower, Troy J. Brinton, and Uday N. Kumar. 2015. *Biodesign: The Process of Innovating Medical Technologies*. 2nd ed. Cambridge University Press.

ARTICLE 10

Cryptocurrencies and resolution 1540 (2004)

Jonathan Brewer

Visiting Professor

King's College London



Abstract

The paper examines the role of cryptocurrencies and their architecture in relation to the proliferation finance measures required by resolution 1540 (2004). The paper analyses resolution 1540's definition of a non-State actor and argues that a cryptocurrency may be considered an entity for the purpose of the resolution. The paper further argues that State-sponsored WMD procurement networks may comprise both State and non-State actors. Cryptocurrencies can be found in case studies of proliferation finance involving both types: Those involving non-State actors may be relevant to the requirements of resolution 1540 (2004). The paper notes increasing evidence of direct sales or purchases of WMD and related materials using cryptocurrencies. Cryptocurrencies (particularly stablecoins) are likely to be increasingly important to proliferation finance in future, and their role should be addressed in the next Comprehensive Review of resolution 1540 (2004), scheduled for 2027.

Keywords

Resolution 1540; proliferation finance; cryptocurrencies; Comprehensive Review 2022.

Author Bio

Dr. Jonathan Brewer OBE is a Visiting Professor at King's College London. He was Acting Coordinator of the UN Panel on DPRK when it was terminated in April 2024, Coordinator of the UN 1540 Group of Experts, and Financial Expert of the UN Panel on Iran. He previously served in the UK Foreign Service (1983 to 2010). He held postings to UK Embassies in Luanda, Mexico City and Moscow, and in London he was seconded to the Cabinet Office as Secretary, Joint Intelligence Committee, and retired from the Service as Head of Counter-Proliferation. His research interests include WMD proliferation, its financing and sanctions. He holds a PhD in Geophysics from Cornell University, New York, NY and a BA in Geology from the University of Oxford, UK.

Introduction

Resolution 1540 (2004) requires States to adopt activity-based measures against proliferation financing. Operative paragraph 2 prohibits the financing of the manufacture, acquisition, possession, development, transport, transfer or use of nuclear, chemical and biological weapons and their means of delivery (hereafter referred to as weapons of mass destruction or WMD) by non-State actors.[1] Operative paragraph 3 (d) requires States to implement controls on financing the export or transshipment of WMD and related materials. The 1540 Committee monitors the progress of States' implementation of measures required under the resolution by means of Committee Matrices which are posted on the Committee website (1540 Committee, n.d.-b).

The Committee conducts a Comprehensive Review of implementation of the resolution every five years, and its report of the most recent Review (in 2022, following a delay caused by the pandemic) recorded an increase of about 5% (to 88%), since the previous review of 2016, of measures taken by States to address the financing requirements of operative paragraph 2. No significant increase was recorded of the measures to control the funding or financing of exports and trans-shipments of WMD and related materials, as required under operative paragraph 3 (d). This remained at the 2016 figure of about 58% (1540 Committee, 2022).

With regard to operative paragraph 2, the Committee recorded that where legislation existed it was related to areas such as terrorism financing, targeted financial sanctions, money-laundering and the role of financial intelligence units, rather than specifically to WMD and related materials. The Committee acknowledged that such legislation could support proliferation finance enforcement efforts with appropriate terms of reference, but it was not the primarily activity-based prohibitions required under resolution 1540 (1540 Committee, 2022, p. 10). With regard to operative paragraph 3 the Committee noted that controls were contained mainly in legislation related to terrorism financing, money-laundering and the establishment of financial intelligence centres, rather than in legislation incorporating specific measures to control the financing of exports and trans-shipments that would contribute to proliferation (1540 Committee, 2022, p. 14). The Committee recorded that more attention was needed, inter alia, on measures to target more effectively the financial flows involving non-State actors engaged in WMD proliferation (1540 Committee, 2022, p. 22). The Recommendations of the Review recorded that the Committee should continue to take note of the use by non-State actors of rapid advances in technology for proliferation purposes (1540 Committee, 2022, p. 25).

Following the Review the Security Council approved Resolution 2663 (2022) (United Nations Security Council, 2022). Under OP12 of this resolution, the Council inter alia, required that the Committee intensify efforts to promote the full implementation of resolution 1540 (2004) by all States "...particularly noting the need for more attention on ... proliferation finance measures". It should be recorded that, despite its use of the term, the Security Council has not agreed a definition of proliferation finance and in its absence this paper uses a definition proposed by a Financial Action Task Force (FATF) Working Group in 2010. FATF regards this as an informal definition but it is appropriate for this discussion because the activities described are all referred to in 1540 (with the addition of brokering and stockpiling). This definition reads:

> "Proliferation financing" refers to the act of providing funds or financial services which are used, in whole or in part, for the manufacture, acquisition, possession, development, export, trans-shipment, brokering, transport, transfer, stockpiling or use of nuclear, chemical or biological weapons and their means of delivery and related materials (including both technologies and dual use goods used for non-legitimate purposes), in contravention of national laws or, where applicable, international obligations (FATF, 2010, p. 11).

Under OP15 of resolution 2663 (2022) the Council also requested the 1540 Committee take note of the continually evolving nature of the risks of proliferation, including the use by non-State actors of rapid advances in science, technology and international commerce for proliferation purposes.

Virtual currencies are one type of developing technology which, in terms of illicit finance in general, is clearly a threat: According to one blockchain analytics company, illicit crypto addresses received $154 billion in 2025 (representing a 162% increase over 2024) (Chainalysis, 2026). According to a separate blockchain analytics company, illicit crypto flows in 2025 amounted to $158 billion (which represented a 145% increase over 2024) (TRM Labs, 2026).

It has long been recognised that in principle virtual currencies including cryptocurrencies can be used by non-State actors to conduct illicit transactions, including to acquire or sell WMD and related materials. The non-State actors considered most likely to be involved in this respect are terrorist groups such as ISIL and Al-Qaida, and others (see for example Analytical Support and Sanctions Monitoring Team, 2024, pp. 19–20).

1. A non-State actor is defined in resolution 1540 (2004) as an individual or entity, not acting under the lawful authority of any State in conducting activities which come within the scope of the resolution.

The potential relevance of virtual currencies to the financial requirements of resolution 1540 (2004) was not included in the agenda of the 2022 Review, and this paper considers the possible impact on these requirements of one particular form of virtual currency, cryptocurrency. The analysis is based on a literature study rather than original research. The paper starts by describing cryptocurrencies and their relevance to PF. It then analyses the definition of non-State actor in 1540 terms (which can include both individuals and entities) and argues that cryptocurrency can be considered an entity. It analyses the role of actors in State-sponsored proliferation networks and argues that such networks may include both State and non-State actors. The paper then considers the use of crypto to circumvent PF requirements of UNSC sanctions related to WMD proliferation, by both State and non-State actors. Although the PF requirements of these sanctions regimes differ in several respects from those of resolution 1540 (2004) the paper concludes that understanding the role of cryptocurrencies in circumventing sanctions may help understand ways in which cryptocurrencies can be used to circumvent 1540's financing requirements, and how States could target 1540-related financial flows more effectively. Finally, the paper makes recommendations for further research on the use of crypto by non-State actors.

The paper is intended to inform the 1540 Committee and UN Member States while they conduct the next Comprehensive Review of 1540, scheduled to take place in 2027 (1540 Committee, n.d.-c).

Background

When resolution 1540 (2004) was approved, it is virtually certain that the references to funding and financing included by the Security Council applied only to fiat currencies. Although the concept of virtual currencies was known in the mid-1990s (see for example The Future of Money: Coins and Banknotes, 2018), only in 2012 did a central bank define the term: "[a] digital representation of value that is neither issued by a central bank or a public authority, nor necessarily attached to a fiat currency but is accepted by natural or legal persons as a means of payment and can be transferred, stored or traded electronically" (European Central Bank, 2012).

Government fiat currencies are centralized and issued by Central Banks, meaning that all transactions made by fiat currencies are recorded on central servers. Customers generally trust such systems because they are based on universally accepted rules and appropriately regulated by national laws and statutes. Financial sector entities are required by most regulatory authorities to conduct consumer due diligence (CDD), which includes know-your-customer (KYC) and anti-money laundering (AML) procedures, and regulatory authorities rely on financial sector entities to implement effective systems and assist law enforcement and security agencies to fight criminals and terrorists. The FATF has incorporated regulatory requirements into its standards, including the need to identify customers and (for entities) beneficial owners, to understand business relationships and to monitor account activity on a risk-based basis (see Recommendations 9-23, FATF, 2026a).

Although the finance and funding requirements of resolution 1540 (2004) were drawn up against that fiat currency background, today, money launderers and terrorist financers increasingly use cryptocurrencies and stablecoins. A cryptocurrency is a type of digital asset. Cryptocurrency networks are based on digital data compiled into blocks ("ledgers"), in accordance with pre-set network rules, and secured by cryptography to protect data integrity and communication security. Crypto is in principle transparent: Every transaction time, date and value is recorded somewhere on the blockchain by means of a "token", a cryptographic representation of the transaction. Users' tokens are held in software called self-custody (or unhosted) "electronic wallets". Wallets have specific addresses on the blockchain which act as pseudonyms for their users. Wallets are also used to store their users' cryptographic keys. Cryptographic keys enable users to conduct crypto transactions and are of two types: A public key, available to anyone in the system, that allows the user to receive cryptocurrency transactions, and a private key (known only to the user) that provides proof of ownership of the public key (Green, n.d.).

Cryptocurrencies are transferred through exchanges. These may be centralised exchanges (CEX), typically run by a commercial enterprise and usually subject to KYC requirements and other regulatory obligations, or decentralised exchanges (DEX) that enable users to trade directly on the blockchain using smart contracts. Smart contracts are pieces of code that automate the actions required of a blockchain transaction, such as borrowing or lending, by self-executing when pre-determined requirements are met, i.e. there is no need for an individual or entity to intervene. DEX trades involve crypto to crypto and typically involve neither KYC nor other due diligence procedures. In contrast to DEXs, peer-to-peer (P2P) exchanges enable buying and selling

cryptocurrencies by directly connecting the people involved without going through a central authority. Some P2P exchanges may impose KYC and additional due diligence procedures. Others may not.

Most cryptocurrency functions are managed by so-called decentralized finance (DeFi) applications based on smart contracts. DeFi operates independently of a central authority that could alter or censor information or requirements. It is potentially more profitable than traditional fiat-based systems in that it enables peer-to-peer transactions without potentially expensive intermediaries. It is also transparent: Smart contracts and transaction records are publicly accessible. But DeFi has downsides: A private user may have few options to recover stolen funds in the absence of intermediaries. And a fundamental issue for crime-prevention authorities is that even though the blockchain is transparent and transactions publicly accessible, linking an individual user to their pseudonym may be difficult if the user has accessed their crypto through an exchange that does not conduct KYC or other form of due diligence.

There are therefore a number of reasons why it could be advantageous for WMD proliferators to switch from fiat currency to cryptocurrency to fund or finance their operations. Crypto companies do not have the same obligations as banks to protect the financial system and identities of wallet owners are not usually monitored. Cryptocurrency transactions can be carried out anonymously: User identities are not always easily linked to crypto addresses and a crypto address may not be linked to a real identity. Illicit transactions can be further hidden using sophisticated electronic laundering tools such as mixers or tumblers (also driven by smart contracts) that pool and mix crypto from multiple users in order to obscure links between a sender of crypto and the recipient and so obscure their transaction history. The lack of centralised regulation reduces opportunities for the authorities to freeze or halt illicit transactions. In addition, criminals using cryptocurrencies can create crypto accounts and digital identities worldwide - they are not restricted by geographical boundaries. Furthermore, tracking and monitoring crypto is technically challenging for most countries, and proliferation financiers can pick countries with weak crypto regulatory systems.

Types of crypto potentially relevant to PF

Cryptocurrencies and their blockchain exist in many different forms with different intended functions (see for example Kraken, 2025; Shammas, 2025). Of particular relevance to 1540's PF measures are forms that could facilitate payments of proliferation-sensitive goods and technologies. Bitcoin, the earliest cryptocurrency created in 2008, is essentially digital cash widely used for payments and so falls into this category. Ether, also widely used for payments, powers Ethereum, a separate blockchain developed to facilitate smart contracts. TRX powers Tron, a derivative of Ethereum focused on scalability and high throughput.

Cryptocurrencies such as Litecoin, Bitcoin Cash (both derivatives of Bitcoin) and Dogecoin were developed to facilitate quick settlements. The Solana blockchain was developed to handle large numbers of transactions.

A problem when using Bitcoin, Ether and other types of crypto for payments is that their value is subject to market demand and can vary unexpectedly. This makes them a potentially unreliable payment mechanism. Stablecoins (digital assets that maintain a stable value against a reference asset) are better in this respect and increasingly widely used. According to FATF, stablecoins represented 30% of all on-chain virtual asset transaction volume in 2025, and 259 types of stablecoins were circulating as of the end of June 2025 (FATF, 2026d). FATF also reported that a blockchain analysis company estimated that stablecoins accounted for 84% of the USD $154B illicit virtual asset transaction volume in 2025, surpassing Bitcoin as the dominant asset for cybercrime-related on-chain transactions (FATF, 2026, p. 12).

Two commonly used stablecoins are USDT which powers the Tether blockchain (Tether, n.d.) and Circle's equivalent stablecoin, USDC (Circle, n.d.). Both are pegged 1:1 to the U.S. dollar. USDT and USDC are backed by reserves held off-chain by their owners (Tether Limited and Circle Internet Group) and they operate via centralized smart contracts that include provisions for the crypto to be frozen or blocked to prevent further transactions, for example if law enforcement finds a connection with criminal activity. Other types of stablecoins may not be frozen so easily. DAI for example is not operated by a single person or company, and its reserves are backed by on-chain collateral, governed by a number of token holders (Eco, 2026). Freezing or blocking DAI is only possible in specific circumstances, for example if a transaction involvesa central exchange.

Furthermore, FATF has identified cases of stablecoins specifically designed to circumvent law enforcement intervention (FATF, 2026b, p. 21).

Privacy coins are an additional form of cryptocurrency potentially relevant to PF. As noted above electronic laundering tools such as mixers or tumblers can be used obfuscate transactions, but anonymity can be further enhanced by using privacy coins such as Monero XMR or Zcash, designed specifically to make tracking difficult (TRM Labs, n.d.-a). For these reasons, and in response to regulatory pressures, privacy coins have been banned by some crypto exchanges.

State and non-State actors

Resolution 1540 (2004) requires States to take measures against non-State actors. The resolution includes a definition of a non-State actor as an "individual or entity, not acting under the lawful authority of any State in conducting activities which come within the scope of this resolution."

But what does this definition mean in practice for law enforcement authorities wishing to formulate or enforce legislation appropriate to non-State actors, or to their converse, State actors?

The definition includes three main elements: “activities”, “lawful authority” and “any state”. The scope of the first element, "activities", is clear: They relate to WMD proliferation and procurement activities as described in the resolution. The second element, “lawful authority” is not further defined but it is reasonable to assume that “lawful authority” and “without lawful authority” can include the presence or absence of a body empowered to extend “lawful authority”, or to the presence or absence of authorization itself (such as a license, or a contract pursuant to such a license).

In the absence of tangible evidence as to whether a contract or license does or does not exist, law enforcement authorities might have to apply more abstract concepts, such as exercise of jurisdiction, power or control. For analytical purposes, the paper suggests that an actor operating under the jurisdiction, power or control of the proliferating State while conducting activities covered by resolution 1540 is a State actor of that State. Other actors conducting such activities would be non-State actors.

Although the third element of the definition, "any State", could in principle be a reference to any one of the 193 UN Member States, it is logical to take it as a reference to the State conducting WMD proliferation. The argument for so doing is as follows, consider a WMD programme in State A: State A wants to use a commercial actor to procure materials from overseas. State A and the commercial actor agree a contract for this purpose. State A has endowed “lawful authority” and so the commercial actor has become a State actor. The commercial actor then procures WMD materials in a separate State, State B. The commercial actor’s activities are illegal in State B. However, the actor is still acting under the lawful authority of State A and is still a State actor of State A. It is neither a State actor of State B nor a non-State actor. The status of the commercial actor's activities with respect to State B or any one of remaining UN Member States is therefore not relevant.

State-sponsored WMD programmes are typically managed by government agencies and departments responsible for planning, funding, procurement and implementation. Materials, equipment or technology not available domestically will have to be procured overseas and the corresponding requirements will be passed onto overseas procurement networks for this purpose. These networks are made up of individuals and entities involved in trading, brokering, shipping, insurance, financing and other commercial concerns (see for example International Institute for Strategic Studies, 2007). For the purposes of this analysis, individuals and entities that are operating with a licence from, or under contract to, government agencies or departments may be considered State actors. But other actors in the network, operating without such licenses or contracts and perhaps far-removed from the proliferating State, may be unaware of a WMD programme and may assume they are conducting standard commercial business. Such actors would fall within the category of non-State actors.

The analysis suggests that State-sponsored WMD proliferation networks may comprise both State and non-State actors. Measures put in place by States to implement resolution 1540 (2004) will apply to those non-State actors.

Is cryptocurrency an entity in 1540 terms?

Although cryptocurrencies are clearly not individuals (i.e. natural persons) they could be considered entities. According to the Cambridge Dictionary the definition of entity includes "something that exists apart from other things, having its own independent existence. The Merriam-Webster Dictionary definition of entity includes "something that has separate and distinct existence and objective or conceptual reality". Either definition applies to cryptocurrencies which possess both independent existence and conceptual reality. The paper argues that cryptocurrencies could be considered entities for analytic purposes under resolution 1540 (2004), although this interpretation is not explicitly addressed in the resolution.

Other elements of the cryptocurrency architecture such as exchanges, DeFi platforms, mixers or other related bodies would also fall under the definition of entity (TRM Labs, n.d.-b).

Examples of crypto in PF

The current study has found no reported examples of crypto being used to violate the financial requirements of resolution 1540 (2004) but case studies exist of crypto being used to violate financial requirements of UNSC sanctions relating to WMD (e.g. Brewer, 2026). Such cases provide clues about the potential use of crypto in violating PF requirements of 1540 although caution is needed: Financial requirements of UNSC sanctions cover PF but they extend to a wide range of requirements outside the scope of 1540 including, but not restricted to, targeted financial sanctions and related asset freezes, sanctions evasions, restrictions or prohibitions on financial services and restrictions or prohibitions on business activities.

Crypto and 1540 - the future

1540 PF requirements revolve around the funding and financing of WMD-related activities, and as noted above the most useful cryptocurrencies will be those designed to facilitate transfers of value or to make payments. Stablecoins, which in principle maintain their value, will be increasingly important in this respect. Stablecoins may be issued by non-State actors like Tether (issuer of USDT) and Circle (USDC), or by State actors (in the case of Central Bank Digital Currencies, backed by a government).

State-linked cybercriminal groups increasingly use stablecoins to launder illicit proceeds (FATF, 2026, pp. 12–13). Laundered funds are converted into stablecoins (usually USDT on the Tron blockchain) and then exchanged into fiat currency through OTC brokers or P2P platforms. According to FATF, the use of stablecoins by threat actors increasingly involves multiple blockchains such as Ethereum, Tron (which is used for moving USDT quickly), Aptos, Avalanche and Solana, connected by tools such as bridges and methods such as cross-chain transactions (FATF, 2026, p. 12).

It must be assumed that 1540-related non-State actors involved in financing and funding WMD proliferation will also try to use stablecoins. Stablecoins will probably be a key element of crypto-related PF in future.

Crypto and 1540 - mitigating the threat

States can implement a number of measures to mitigate the threat of non-State actors using cryptocurrencies to fund or finance WMD proliferation. Of fundamental importance is ensuring that effective national systems are in place to identify, track and monitor the related crypto transactions. This will be challenging to most States that possess neither expertise nor resources, and such States could seek support through the 1540 Committee's assistance mechanism (1540 Committee, n.d.-a) or directly from States that are better equipped and experienced. Alternatively, States could contract blockchain analytics companies. Those companies would need clear guidance on how to distinguish between non-State and State actors.

On-chain transactions cannot be stopped in the same way that, for example, a correspondent bank can block a cross-border financial transaction. Because blockchain architecture is decentralised, regulatory authorities need to consider how to monitor, control and if necessary target the blockchain as a whole (Keating & Redslowska, 2026). In particular they need to focus on the workings of the crypto exchanges which, because they govern the interaction of users of the chain with other users and with other chains, are key to the

functioning of the system.

If resources are insufficient to allow either the implementation of effective national systems or employment of blockchain analytics companies, States should implement systems to monitor any crypo/fiat currency exchanges that are operating in their jurisdictions. Unless such exchanges are centralised they are likely to be non-State actors and, if transacting business related to proliferation of WMD or related materials, likely to be operating in violation of 1540's PF measures. States should identify individuals or entities offering cashing-out services by means of such exchanges (Brewer, 2026).

1540, crypto and FATF standards

FATF standards related to PF are limited to UNSC targeted financial sanction requirements and to national assessments of risk based on those requirements but do not cover resolution 1540 (2004). Despite the narrow focus of its standards however, FATF encourages States to consider information relating to broader issues of counter-proliferation (FATF, 2021, p. 4), and has published a number of guidance documents to help States to monitor and control crypto activity. These are potentially helpful for monitoring and controlling 1540 PF activity, and include a list of 69 jurisdictions which currently account for approximately 97% of the global virtual asset market (FATF, 2026b, p. 27), and an analysis of risks associated with offshore VASPs (FATF, 2026b).

FATF Member States should also implement Recommendation 15 which requires States to apply counter-PF measures to virtual assets and VASPs. Although like other FATF Recommendations it does not cover resolution 1540 (2004), the act of implementation may help deter 1540 PF crypto activities. The level of implementation of this Recommendation is low: In 2024 75% of jurisdictions were partially or not compliant (FATF, 2024).

New developments

Recent evidence suggests that crypto is being used to purchase or sell proliferation-sensitive goods and materials directly, without first converting it to fiat currencies (FATF, 2026, p. 3).

The benefits of conducting proliferation-sensitive transactions directly in crypto are obvious: There is no need to have a reliable cashing-out mechanism in place, potentially vulnerable to law enforcement action, prior to procurement. Furthermore (and although not relevant to implementation of resolution 1540 (2004)) crypto offers a potential channel to circumvent financial sanctions. It is likely that non-State actors involved in financing and funding WMD proliferation may increasingly focus on conducting transactions in crypto rather than fiat currency in future.

Conclusions

Resolution 2663 (2022), approved following the 2022 Comprehensive Review of resolution 1540 (2004), requests that the 1540 Committee take note of the use by non-State actors of rapid advances in technology for proliferation purposes. Cryptocurrencies, which were not addressed by the Committee during the Review, are a clear example of a rapid advance in technology.

The Committee's report of the Review records that States need to pay more attention to measures to target more effectively the financial flows involving non-State actors engaged in WMD proliferation. Such targeting should clearly extend to crypto: This paper explores whether cryptocurrencies could be regarded as entities for the purpose of resolutoin1540.

Cryptocurrencies and crypto architecture that facilitate payments or other types of value transfer are particularly relevant to 1540 and PF. They currently include, but are not restricted to, Bitcoin and Tether, and the Ethereum and Tron blockchains. Stable coins, designed to maintain their value against fiat currencies, are increasingly used for illicit finance and may become a common feature of PF.

Cryptocurrencies are increasingly used for direct sales or purchase of WMD and related materials. Such sales avoid the need for cashing-out mechanisms, potentially vulnerable to law enforcement actions. They too may become a common feature of PF

These are all issues that should be addressed during the next 1540 Comprehensive Review, scheduled for 2027. In particular, the Review may wish to examine ways to strengthen the requirements of resolution 1540 (2004) in the light of the PF threat from crypto.

List of citations

1540 Committee. (n.d.-a). 1540 Assistance Mechanism. United Nations. Retrieved July 26, 2026, from https://www.un.org/en/sc/1540/assistance/general-information.shtml

1540 Committee. (n.d.-b). 1540 Committee approved Matrices [Dataset]. Retrieved un.org/en/sc/1540/national-implementation/1540-matrices/committee-approved-matrices.shtml

1540 Committee. (n.d.-c). 2027 Comprehensive Review. United Nations Security Council. Retrieved un.org/en/sc/1540/comprehensive-and-annual-reviews/2027-comprehensive-review.shtml

1540 Committee. (2022). Final document on the 2022 Comprehensive Review of the status of implementation of resolution 1540 (2004) (S/2022/899). United Nations Security Council. https://www.un.org/en/sc/1540/comprehensive-and-annual-reviews/2022-comprehensive-review.shtml

Analytical Support and Sanctions Monitoring Team. (2024). Thirty-fourth report of the Analytical Support and Sanctions Monitoring Team submitted pursuant to resolution 2734 (2024) concerning ISIL (Da'esh), Al-Qaida and associated individuals and entities (S/2024/556). United Nations Security Council. https://docs.un.org/en/S/2024/556

Brewer, J. (2026). A Focus On "Cashing Out": One Way to Combat DPRK's Crypto Thefts. 38 North, Stimson Center. https://www.38north.org/2026/03/a-focus-on-cashing-out-one-way-to-combat-dprks-crypto-thefts/

Chainalysis. (2026). The Chainalysis 2026 Crypto Crime Report. https://www.chainalysis.com/reports/crypto-crime-2026/

Circle. (n.d.). The full-stack platform for the internet financial system. Retrieved July 22, 2026, from https://www.circle.com/

Eco. (2026, July 16). What Is DAI? Decentralized Stablecoin Explained 2026 | Support. https://eco.com/support/en/articles/15210329-what-is-dai-decentralized-stablecoin-explained-2026

European Central Bank. (2012). Virtual currency schemes. European Central Bank. https://doi.org/10.2866/47380

FATF. (2010). Combatting Proliferation Financing: A Status Report on Policy Development and Consultation. Financial Action Task Force. https://www.fatf-gafi.org/en/publications/Financingofproliferation/Combattingproliferationfinancingastatusreportonpolicydevelopmentandconsultation.html

FATF. (2021). Guidance on Proliferation Financing Risk Assessment and Mitigation. https://www.fatf-gafi.org/en/publications/Financingofproliferation/Proliferation-financing-risk-assessment-mitigation.html

FATF. (2024). Targeted Update on Implementation of the FATF Standards on Virtual Assets and Virtual Asset Service Providers. Financial Action Task Force. https://www.fatf-gafi.org/content/fatf-gafi/en/publications/Fatfrecommendations/targeted-update-virtual-assets-vasps-2024.html

FATF. (2026). Targeted Report on Stablecoins and Unhosted Wallets Peer-to-Peer Transactions. Financial Action Task Force. https://www.fatf-gafi.org/en/publications/Virtualassets/targeted-report-stablecoins-unhosted-wallets.html

FATF. (2026a, June). The FATF Recommendations. Financial Action Task Force. fatf-gafi.org/en/publications/Fatfrecommendations/Fatf-recommendations.html

FATF. (2026b, July). Seventh Targeted Update on Implementation of the FATF Standards on Virtual Assets/VASPs. https://www.fatf-gafi.org/en/publications/Fatfrecommendations/targeted-updated-virtualassets-vasps-2026.html

Green, G. (n.d.). Research Guides: Fintech: Financial Technology Research Guide: Cryptocurrency & Blockchain Technology [Research guide]. Retrieved July 24, 2026, from https://guides.loc.gov/fintech/21st-century/cryptocurrency-blockchain

International Institute for Strategic Studies (Ed.). (2007). Nuclear black markets: Pakistan, A.Q. Khan and the rise of proliferation networks ; a net assessment. IISS.

Keating, T., & Redslowska, K. (2026). Enhancing the EU Response to Crypto-Enabled Sanctions Circumvention. Royal United Services Institute. https://www.rusi.orghttps://www.rusi.org

Kraken. (2025). Https://www.kraken.com/learn/types-of-cryptocurrency. https://www.kraken.com/learn/types-of-cryptocurrency

Shammas, S. (2025, February 26). 9 Cryptocurrencies That Make Payments Easy and Efficient. Cryptonix. https://www.cryptonix.com/blog/9-cryptocurrencies-that-make-payments-easy-and-efficient

Tether. (n.d.). What are Tether tokens and how do they work? Retrieved July 22, 2026, from https://tether.to/ru/how-it-works/

The Future of Money: Coins and Banknotes (2018). https://financialservices.house.gov/uploadedfiles/115-113.pdf

TRM Labs. (n.d.-a). Privacy Coins. TRM Labs. TRM Glossary. Retrieved https://www.trmlabs.com/glossary/privacy-coins#riskbased-approach-1

TRM Labs. (n.d.-b). Why Entity Due Diligence Matters in Crypto: Risks, Compliance & Trust. Retrieved July 22, 2026, from https://www.trmlabs.com/glossary/entity-due-diligence

TRM Labs. (2026). 2026 Crypto Crime Report. https://www.trmlabs.com/reports-and-whitepapers/2026-crypto-crime-report

United Nations Security Council. (2022, December 1). S/Res/2663 (2022). https://docs.un.org/en/S/RES/2663%20(2022)

ARTICLE 11

Making Emerging Risks Visible under UNSCR 1540: Academic Knowledge Networks and the Conditions for Distributed Anticipatory Infrastructure

Konstantinos Papadikis
Associate Vice President for Education
Xi'an Jiaotong-Liverpool University



Abstract

United Nations Security Council resolution 1540 (2004) obliges States to prevent non-State actors from acquiring nuclear, chemical and biological weapons, their means of delivery, and related materials, equipment and technology. Every control measure it requires presupposes that the items and knowledge concerned have first been recognised as being of concern, and convergence across the life sciences, computational methods, laboratory automation and advanced manufacturing has made that recognition a judgement falling between specialisms. Building on scholarship in anticipatory governance, this article develops a four-process framework for assessing when heterogeneous academic and treaty-facing arrangements may contribute to a distributed anticipatory infrastructure: horizon sensing, knowledge translation, awareness diffusion and institutional learning. It pairs the framework with an evidentiary test that treats a connection as established only where documentation traces the output of one process into the operation or feedback of another, and applies both to five purposively selected cases examined through public records. The component capacities are present and are carried by arrangements of different kinds. Partial connections meeting the test are traceable within individual cases, but no case documents a configuration connecting all four. Academic knowledge networks therefore form one component of such an

infrastructure rather than constituting one in themselves, and they contribute upstream of obligations that rest on States. The contribution is analytical rather than empirical: co-existing capacities do not by themselves establish an infrastructure, and the article specifies what evidence a claim of that kind would require.

Keywords: UNSCR 1540; non-State actors; dual-use technology; research security; academic knowledge networks; anticipatory awareness

1. Introduction

United Nations Security Council resolution 1540 (2004) requires States to prevent non-State actors from acquiring nuclear, chemical and biological weapons, their means of delivery, and the materials, equipment and technology associated with them. The obligations it creates are legal in character and are addressed to States. What they presuppose is not.

Every measure the resolution requires, from accounting for and securing sensitive items to physical protection, border controls and export and trans-shipment controls, operates on items, transactions and bodies of knowledge that must first be recognised as being of concern. Where the relevant technologies are long-established, recognition is more institutionally stabilised, even if classification and end-use judgements remain contested, and the principal difficulty lies in implementation and enforcement. Where they are not, the position reverses. Developments in the life sciences, in computational and data-driven methods, in laboratory automation and in advanced manufacturing increasingly draw on one another, and a result in one field may acquire significance only when read against a capability in another. The judgement that a development matters for the resolution's purposes then falls between specialisms, and it is not obvious who is placed to make it.

This article examines the arrangements through which such judgements are formed. Its central question is how academic knowledge networks, the durable relationships among researchers, universities, professional societies, scientific advisory bodies and international organisations through which specialised knowledge is produced, interpreted and retained, contribute to anticipatory awareness of emerging and converging technological developments relevant to resolution 1540 (2004). Two subsidiary questions follow. What processes do such arrangements carry? And what limits their capacity to identify and communicate dual-use concerns?

The answer developed here is deliberately narrow. Building on scholarship in anticipatory governance, the article develops a four-process framework for assessing when heterogeneous academic and treaty-facing arrangements may contribute to a distributed anticipatory infrastructure relevant to resolution 1540 (2004): horizon sensing, knowledge translation, awareness diffusion and institutional learning. The first three are adaptations of capacities identified in that literature; the fourth is an article-specific extension addressed to retention and feedback. The framework is paired with an evidentiary test for connection, since without one the question of whether processes are joined is settled by intuition. Applied to five documented cases, the framework yields a result that constrains the article's own claim. The component capacities are present, carried by arrangements of different kinds: a standing academic network, a treaty-facing advisory interface, a project coalition, a continuing State-academia outreach process, and a single outreach event. Partial connections meeting the test are traceable, but no case exhibits a configuration connecting all four processes. Academic knowledge networks therefore contribute conditionally rather than constituting an infrastructure on their own. That is a smaller claim than the framework might invite, and the evidence assembled here does not support a larger one. The contribution is accordingly analytical rather than empirical: co-existing capacities do not establish an infrastructure, and this article specifies what would.

Three boundaries should be stated at the outset. This is a conceptual article rather than an evaluation of effectiveness, as Section 3.4 sets out. The obligations under the resolution rest on States throughout, and nothing described here transfers or substitutes for them; where the resolutions engage academic expertise, they do so in support of implementation and, when the 1540 Committee draws upon it, with the consent of the State concerned. And the technologies discussed appear as broad categories, since the argument concerns how significance is recognised rather than what any technique can do.

The article proceeds as follows. Section 2 sets out the resolution's scope and traces how its texts have treated knowledge, intangible transfer and academic participation between 2004 and 2022. Section 3 characterises the recognition problem, positions the framework in relation to the anticipatory governance literature, defines its two central concepts and states the approach to case selection. Section 4 develops the four processes and examines the cases. Section 5 probes the framework across four broad technology domains. Section 6 sets out five tensions that constrain the arrangements described, and Section 7 draws out the implications for awareness and implementation-support discussions.

2. UNSCR 1540, dual-use knowledge and technological convergence

2.1 The scope of the resolution

Resolution 1540 (2004) was adopted under Chapter VII of the Charter on 28 April 2004. Its operative core is short. States shall not support non-State actors seeking nuclear, chemical or biological weapons and their means of delivery; States shall adopt and enforce laws prohibiting non-State actors from developing, acquiring, possessing, transporting, transferring or using such weapons; and States shall establish domestic controls, including accounting and securing measures, physical protection, border controls and export and trans-shipment controls, over what the resolution calls related materials (United Nations Security Council [UNSC], 2004, paras. 1–3). The resolution defines that last term itself, and the definition repays attention. Related materials are "materials, equipment and technology" covered by relevant treaties and arrangements or included on national control lists, which could be used for the design, development, production or use of the weapons and delivery systems concerned (UNSC, 2004, footnote). Technology sits inside the definition from the beginning; the resolution's object was never only physical stock.

The resolution also addresses engagement beyond government. States are called upon to develop appropriate ways to work with and inform industry and the public regarding their obligations under such laws (UNSC, 2004, para. 8(d)). Although academia is not named in the 2004 text, later resolutions place it explicitly within civil-society outreach (UNSC, 2016, preamble and para. 29; UNSC, 2022, para. 25).

2.2 From physical transfer to knowledge flow

Subsequent resolutions made the knowledge dimension progressively more explicit. In 2011, resolution 1977 emphasised the need for States to strengthen export controls and "to control access to intangible transfers of technology and to information that could be used for weapons of mass destruction and their means of delivery", and described full implementation as a long-term task requiring continuous effort (UNSC, 2011, preamble). The same resolution encouraged the 1540 Committee, in preparing technical reference material, to draw on relevant expertise including civil society and the private sector, "with, as appropriate, their State's consent" (UNSC, 2011, para. 12), and invited the Committee to develop ways of utilising and maintaining expertise, including that of former members of its group of experts (UNSC, 2011, para. 22(d)).

Resolution 2325 (2016) carried both threads further. It repeated the intangible-transfer concern as an operative encouragement (UNSC, 2016, para. 13); it called upon States to take into account "the evolving nature of risk of proliferation and rapid advances in science and technology" in implementation, and requested the Committee to do likewise in its own work (UNSC, 2016, paras. 7–8); and it acknowledged in its preamble the positive role that civil society, "inter alia industry and academia", could play in effective implementation, including by raising awareness (UNSC, 2016, preamble). Academic communities entered the operative text in the same resolution,

through the direction that outreach events include representatives of civil society, industry and academia (UNSC, 2016, para. 29) and the encouragement that the Committee draw on relevant expertise "including industry, scientific and academic communities, with, as appropriate, their States' consent" (UNSC, 2016, para. 30).

Resolution 2663 (2022) extended the Committee's mandate to 2032 and made the surrounding machinery routine: comprehensive reviews on a fixed cycle, an annual programme of work and review, technical reference guides that are developed and revised, a periodically reviewed assistance-matching mechanism, feedback from both recipients and providers of assistance, regular open briefings, and continued outreach including academia (UNSC, 2022, paras. 2–5, 13, 17, 20, 25–27). The arc from 2004 to 2022 is not one in which an epistemic dimension was absent and then discovered, but one in which a dimension present in the original text, namely technology within related materials and obligations that must be communicated to those who handle them, was progressively specified and then routinised.

Two features of that development matter for what follows. The first is the consent formula: the 2011 and 2016 resolutions condition the Committee's drawing on external expertise, as appropriate, on the consent of the State concerned, and describe such expertise as assisting States in their implementation. What the texts establish, in that setting, is a state-centred, consent-conditioned supporting role. The second is the recurrence of review and feedback as design features of the machinery itself. Whether academic networks exhibit anything comparable is a separate question, and one the texts cannot answer.

2.3 Convergence and the problem of recognition

The technological environment those texts address has changed in a way that bears on recognition. Developments in the life sciences, in computational and data-driven methods, in laboratory automation and in advanced manufacturing increasingly draw on one another, so that a result in one field acquires significance only when read against a capability in another. Convergence of this kind does not merely add to the volume of material to be assessed; it changes who is competent to assess it. A specialist may be entirely current in their own field and still unable to say whether a result matters for the objectives of resolution 1540, because the significance lies in a combination no single field owns.

This is a problem about knowledge before it is a problem about control. Control lists, border measures and export controls operate in part on categories that must first be recognised and described. Control systems are not helpless before novelty: catch-all provisions, end-use and end-user controls and licensing judgements are intended to reach items and transactions that no list anticipated. They can mitigate the recognition problem without removing it, since each still requires someone to judge that a particular item or transaction is of concern, and that judgement depends on the same distributed expertise.

The connection to the resolution is direct. Recognising dual-use significance earlier brings a development within reach of the awareness-raising, education and knowledge-exchange activity that the resolutions associate with academia and industry (UNSC, 2016, preamble, paras. 29–30); and that activity bears on the resolution's central objective, since the materials, equipment and technology States must control have first to be understood as being of concern.

3. The anticipatory awareness gap

3.1 Four difficulties of early recognition

Four difficulties are distinguished here for analytical purposes; they are not offered as an exhaustive taxonomy, and they overlap in practice. What they share is that better rules alone do not resolve them, because each concerns epistemic conditions that rules and controls cannot create by themselves.

The first is that development outruns vocabulary. A capability can exist for some time before the research community settles on a name for it, and control frameworks work with named categories. During the interval, the thing is describable only by those close enough to the work to describe it without a term.

The second is that significance is distributed across fields. What one community sees as a routine methodological advance may carry implications visible only to another. Neither is negligent; each sees accurately within its own frame. The judgement that matters exists in neither place, and comes into being only if the frames are brought together.

The third is that dual-use boundaries are genuinely uncertain rather than merely disputed. The same research may hold medical, industrial, scientific and security implications at once, and the balance among them can shift with context and subsequent work. Treating such boundaries as settled misdescribes the object; treating them as unknowable abandons the assessment.

The fourth is that knowledge moves in ways physical measures were not built to follow. Software, models, data and remote collaboration do not cross borders as consignments do. The Council has repeatedly directed attention to intangible transfers of technology and to information (UNSC, 2011, preamble; UNSC, 2016, para. 13), which acknowledges the incompleteness of the physical template rather than solving the problem it names.

Together these describe a gap between the pace of technological change and the capacity of the relevant communities to interpret it. This is the gap the article calls one of anticipatory awareness.

3.2 Theoretical position

The framework developed below draws on anticipatory governance, in which Barben et al. (2008) identified foresight, engagement and integration as the constitutive capacities, and which Guston (2014) elaborated as an ensemble including long-term reflective capacity building rather than one-off exercises. Two adaptations carry that work into the present context: the relevant communities are narrower than the publics addressed in the nanotechnology literature where the concept was developed, and the institutional setting is shaped by a Chapter VII resolution that places participation within a state-centred, consent-conditioned frame the original formulation did not have to address. The framework therefore adapts these capacities with modified meanings and adds a fourth process addressed to retention and feedback. It is not a one-to-one relabelling: the four processes are reconstructed for the narrower problem examined here. Two further bodies of work supply the remaining scaffolding: the knowledge-networks literature, for a vocabulary that describes such arrangements without assuming they are organisations (Phelps et al., 2012), and the organisational learning literature, for the movement by which what individuals and groups know may become institutionalised and available to a collective over time (Crossan et al., 1999).

3.3 Two core concepts

Academic knowledge networks

In this article, academic knowledge networks are understood as relatively durable patterns of relationship among researchers, universities, professional societies, scientific advisory bodies and international organisations, through which specialised knowledge about science and technology is produced, interpreted, reviewed and retained. The formulation follows that literature in treating social relationships as the medium through which knowledge is created, circulated and taken up (Phelps et al., 2012). Two criteria separate a network from an audience. First, academic actors play a structural role in producing, translating or reviewing knowledge; attendance at an event as invited listeners does not qualify. Second, the relationships persist across institutions and across cycles of activity, since a single workshop, however well attended, is not a network. Both are kept separate from the question the framework later asks about retention: whether a network also retains and feeds back what it learns is evidence about one of the four processes, not a condition of being a network at all. Holding them apart is what makes failure at the fourth process observable rather than defined away. The category is adjacent to, but not the same as, an epistemic community in the sense of Haas (1992): it does not require that participants share a single set of normative and causal beliefs, only that they are structurally engaged in the same knowledge work.

Anticipatory awareness

Anticipatory awareness refers here to a composite epistemic condition: one in which technical and non-proliferation communities are able to recognise, interpret and communicate the dual-use significance of emerging scientific and technological developments before or alongside their recognition in formal guidance, control lists or documented incidents. The term is anticipatory in the sense developed in the governance literature rather than a predictive one (Guston, 2014): the condition is built in advance of need, and nothing in the framework predicts where proliferation attempts will occur. What it describes is whether the communities concerned possess the concepts, relationships and retained experience needed to see the significance of technological change early enough for awareness, education and knowledge exchange to bear on it.

3.4 A note on approach and case selection

This is a conceptual, framework-building article, not an evaluation of effectiveness. It does not measure whether any arrangement has reduced the probability of a proliferation event, and no claim of that kind should be read into the analysis. The cases were selected purposively, not sampled. Three criteria governed selection: the availability of sufficient public documentation, from primary and academic sources; variation in institutional setting, carrier type and centre of knowledge production, so that the framework is not tested only against arrangements built to resemble it; and the capacity of each case to test a boundary or failure mode, not merely to illustrate the framework's operation. No requirement was imposed that a case exhibit a minimum number of processes; such a requirement would have selected for the conclusion, and in the event the cases differ sharply in how many processes they carry and in whether the links between them can be traced at all.

Selection on these grounds has a known limitation: cases chosen for their documentary visibility may over-represent well-resourced institutions, and purposive selection cannot rule out confirmation bias. Two safeguards are applied. Complete cycles are treated as findings to be demonstrated, never as conditions of entry. One case retained here falls outside the category of network altogether, and is kept for the boundary it marks. And the

language used for each case is graded by its relationship to the resolution, as directly connected, closely related to its objectives, or analogous, so that illustrative material is not presented as implementation practice. The analytical contribution lies in the framework; the cases illustrate and probe it.

4. Four processes and the conditions for distributed anticipatory infrastructure

4.1 Four processes, one cycle

Four processes follow from the adaptation set out in Section 3.2: horizon sensing, knowledge translation, awareness diffusion and institutional learning. The fourth is an article-specific extension foregrounding retention and feedback. The processes are analytically distinct but empirically interdependent, and inherit the ensemble logic of the original formulation rather than a fixed order. A training course with no retained record is not yet infrastructure; a scanning exercise whose findings never cross disciplinary boundaries is not yet awareness.

One point governs how the cases are read. These are processes, not properties of any single type of body, and the four together state a condition rather than an achievement: an arrangement, or set of arrangements, amounts to anticipatory infrastructure for the purposes of resolution 1540 (2004) only where the processes are both present and connected. Nothing requires that one arrangement carry all four.

Connection needs a stated test, since without one the judgement rests on intuition. A connection is claimed here only where the documentary record traces an observable output of one process into the uptake, operation or feedback mechanism of another. Institutional co-presence, compatible mandates and chronological succession are not sufficient. That test is deliberately demanding, and applying it has a consequence for what the cases can show, which Section 3.4 has already flagged and Section 4.7 returns to: cases selected for variation in carrier type are not thereby a sample capable of tracing one configuration end to end.

4.2 Horizon sensing

Horizon sensing is the earliest identification, within research communities, of scientific and technological developments whose dual-use significance is not yet settled. Its carriers are ordinary organisational nodes and communicative venues: journals, conferences, professional societies, scientific advisory bodies. It is not intelligence collection, and it is not formal technology assessment by governments; the evidence examined here is publicly documented. The observable outputs are horizon-scanning reports, advisory-body assessments and review procedures that flag dual-use concerns. The Security Council has itself made the object of this process explicit: States are called upon to take into account "the evolving nature of risk of proliferation and rapid advances in science and technology" in their implementation of the resolution (UNSC, 2016, para. 7), and the 1540 Committee is requested to take note of the same developments in its own work (UNSC, 2016, para. 8). Sensing fails in two characteristic ways: fields without security-aware reviewers miss signals, and communities that flag too much erode their own credibility.

4.3 Knowledge translation

Knowledge translation converts technical developments into terms that non-proliferation, legal and research-security communities can assess, and carries their concerns back into technical language. It is a two-way process, and it is analytically distinct from, and does not itself entail, policy advocacy. Translation typically leaves traces: joint publications across disciplinary boundaries, advisory reports prepared for treaty bodies, shared terminology that persists in later documents. The textual basis for this process in the 1540 context is the Council's encouragement of the Committee to draw on relevant expertise, "including industry, scientific and academic communities, with, as appropriate, their States' consent" (UNSC, 2016, para. 30). That formulation is precise about roles, and this article follows its precision: what the resolution supports is a state-centred, consent-conditioned supporting role for academic expertise. The characterisation of that role as upstream and epistemic is the analytical position of this article, argued through the framework and cases, not a conclusion of the text itself. Translation fails where expertise remains fragmented, that is, where technical and security communities each hold half of a dual-use judgement and no mechanism joins them.

4.4 Awareness diffusion

Awareness diffusion spreads translated understanding through education, training, publication and professional-society activity. It is voluntary and educational, not compliance enforcement. Its outputs are visible in curricula, codes of conduct and training records. This is the process the resolutions came to recognise across the sequence traced in Section 2, from an obligation to inform industry and the public in 2004 to the naming of academia and the direction of outreach towards it in 2016 (UNSC, 2016, preamble, para. 29). Diffusion has a characteristic failure mode: material reaches audiences already persuaded of its relevance, while awareness fades in institutions that lack carriers to retain it.

4.5 Institutional learning

Institutional learning is the retention of cases, terminology, review experience and relationships within institutions and networks, such that awareness survives the departure of individuals and feeds back into the next round of sensing. The organisational learning literature describes this as the movement by which what individuals and groups come to know becomes institutionalised and available to the collective over time (Crossan et al., 1999); the process named here is the analogue of that movement within networks and institutional interfaces rather than within a single firm. It is not legislative institutionalisation; the mechanisms are internal to universities, societies and organisations. Observable evidence includes iterated review mechanisms, successive editions of guidance, arrangements for retaining expertise across personnel changes, and standing expert networks. The 1540 machinery offers an institutional analogue in its own review-and-feedback design (UNSC, 2022, paras. 3, 13, 17, 20), but the analogy is a mirror rather than proof: that the Committee's machinery is built for retention does not demonstrate that academic networks learn, and the cases below carry the burden of answering that. Learning fails through memory loss, meaning staff turnover without handover, and through ritual: reviews that recur on schedule while feeding nothing back.

4.6 Table 1 and the cycle

Table 1. Four processes contributing to distributed anticipatory infrastructure relevant to UNSCR 1540

Process	What it does	Observable evidence	Relevance to UNSCR 1540	Failure mode
Horizon sensing	Identifies the dual-use significance of emerging developments within research communities	Scanning reports; advisory assessments; flagged concerns	Evolving-risk provisions (UNSC, 2016, paras. 7–8)	Missed signals where review is not security-aware; credibility loss through over-flagging; restrictions that narrow expert scrutiny
Knowledge translation	Renders technical developments into non-proliferation terms, and the reverse	Joint publications; advisory reports; shared terminology	Consent-conditioned use of academic expertise (UNSC, 2016, para. 30)	Fragmented expertise; meaning distorted across disciplinary boundaries
Awareness diffusion	Spreads translated understanding through education, training and publication	Curricula; codes of conduct; training records	Awareness-raising role of academia (UNSC, 2016, preamble, para. 29)	Reach limited to persuaded audiences; unequal access; limited adaptation to local conditions
Institutional learning	Retains cases, terminology and review experience, and feeds them back into earlier processes	Revised guidance; iterated review mechanisms; standing expert networks	Structural analogue in review-and-feedback design (UNSC, 2022, paras. 3, 13, 17, 20)	Memory loss through turnover; ritualised review without feedback; resource gaps preventing retention

Source: Author's analytical framework, adapting the anticipatory governance capacities of Barben et al. (2008).

Figure 1. Distributed processes of anticipatory awareness

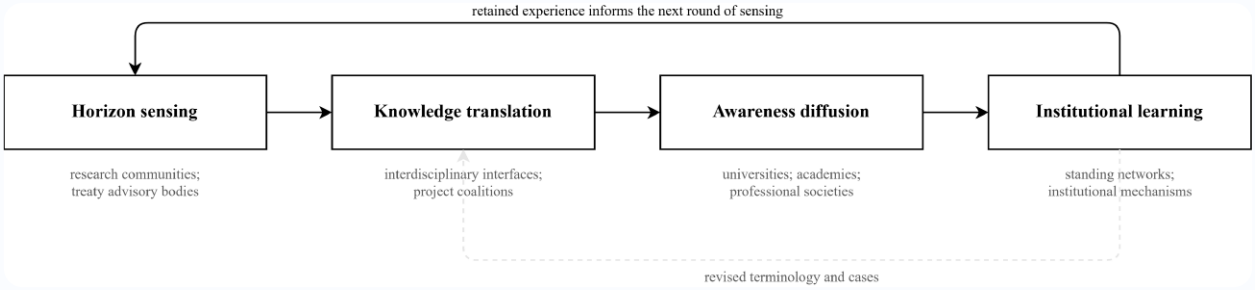

Note. The four processes may contribute to anticipatory awareness relevant to UNSCR 1540, with indicative carriers shown beneath each. The figure represents distributed processes, not the internal structure of a single organisation. The cycle is heuristic rather than a necessary sequence, and observable feedback is the feature through which repeated activity may become infrastructural. Source: Author's analytical framework.

4.7 Anchor cases: how the framework meets the record

Five cases probe the framework against publicly documented practice. They are deliberately of different kinds, and are taken in this order: a single outreach episode that falls outside the category altogether, a continuing State-academia outreach process, the interface at which academic communities meet treaty machinery, a standing academic network, and a project coalition on a longer line of work. That heterogeneity is the substance of what they show. Each is graded by its relationship to resolution 1540 (2004), and the grading records the closeness of that relationship rather than any judgement of effectiveness.

A single outreach seminar: the boundary of the category. The Council has directed the 1540 Committee to organise outreach events including representatives of civil society, industry and academia (UNSC, 2016, para. 29; UNSC, 2022, para. 25). A seminar at the Faculty of Law of the University of Buenos Aires on 20–21 September 2022, organised with the Inter-American Committee against Terrorism of the Organization of American States, brought the resolution's obligations before postgraduate students and national officials, with the stated objective of strengthening their knowledge of it (1540 Committee, 2022). The case is directly connected to the resolution, and dissemination activity is plainly present. What it does not establish marks the category's boundary: a single seminar is an outreach episode, not an academic knowledge network, since relationships persisting across institutions and cycles cannot be met by one event, and the material examined here records no retention on the academic side.

A continuing State-academia outreach process: activity that reaches a repository. The same textual basis supports a more sustained arrangement. In 2023 the German Federal Foreign Office, supported by the United Nations Office for Disarmament Affairs, the German Federal Office for Economic Affairs and Export Control and the Fraunhofer Society, launched the Erlangen Initiative, an outreach process directed at academia and modelled on the Wiesbaden Process, which has addressed industry since 2012. A second conference followed in 2024, and its discussions were compiled into a document, *Export Controls: The Importance of Engaging Academia*, transmitted to the Chair of the 1540 Committee by note verbale on 7 July 2025, expressly "in accordance with operative paragraph 8 (d) of UNSCR 1540 (2004)", with the request that it be uploaded to the Committee's website, potentially within its section on experiences shared, lessons learned and effective practices (Permanent Mission of the Federal Republic of Germany to the United Nations, 2025). The Committee's index of that section lists the submission under Germany, dated 7 July 2025 (1540 Committee, n.d.).

This is the clearest instance in the present set of an output travelling from one process towards another: deliberative activity produced a codified document, and that document entered a standing repository maintained by the Committee. Three limits keep the reading exact. The record documents outreach and dissemination activity, not any measured change in what audiences know or do, so it cannot be read as evidence that awareness was raised. Deposit is entry into a retention channel rather than proof that retention has occurred: the material examined here records no subsequent uptake, reuse or feedback. And a conference series with an accompanying document is not, on this evidence, a network in the sense of Section 3, since continuity of process is not continuity of relationships. The case is directly connected to the resolution.

Advisory boards at a treaty interface: complementary capacities. The Scientific Advisory Board of the Organisation for the Prohibition of Chemical Weapons (OPCW) reviews developments in science and technology and reports them in terms the treaty community can act upon; its temporary working group on the convergence of chemistry and biology, whose final report the Board endorsed in 2014, translated a diffuse research trend into an assessed treaty-relevant concern (OPCW, 2014). A second standing body, the Advisory Board on Education and Outreach, advises the Director-General on education, outreach and awareness-raising (OPCW, n.d.). The record establishes the co-presence, within one treaty interface, of sensing and translation capacities alongside a dedicated education capacity. It does not establish a causal passage from any particular assessment to any particular educational product: the documentary record contains no such tracing, and this article does not assert one. The

distinction matters for the framework, since co-presence is easily mistaken for connection. The OPCW is also an international organisation rather than an academic network, so the case documents an interface. It is closely related to the objectives of the resolution, which directs States towards cooperation within the OPCW framework (UNSC, 2004, para. 8(c)).

A nuclear-security education network with observable retention. The International Nuclear Security Education Network (INSEN), established in 2010 under the auspices of the International Atomic Energy Agency (IAEA) to support, sustain and promote nuclear security education, is the case in which the later processes can be traced furthest. What it does not document is horizon sensing in the sense used here: the network's attention is directed at educational provision rather than at emerging technological developments, and nothing in the record shows it identifying a dual-use signal. Its translation evidence is weaker than it may appear, documenting the development of educational materials against an agency curriculum rather than the rendering of technological developments into non-proliferation terms. The case is evidence about diffusion and retention, and about the link between them, not about the cycle entire. By 2023 the network comprised 204 member institutions from 72 IAEA member States and thirteen observers, with the largest concentrations in Europe, Asia and Africa (Metwally et al., 2024). Both criteria of the network definition are met: academic actors hold structural roles in producing and reviewing material rather than receiving it, and relationships persist across institutions and annual cycles through three standing working groups covering educational materials and methodologies; programme, curriculum and faculty development; and knowledge management and promotion (Metwally et al., 2024).

The fourth process is visible in two mechanisms rather than inferred from the network's longevity. The first is codification: members of the first working group supported a technical document on sustaining and assessing the impact of nuclear security education programmes, intended, in the network's own description, to help current and future members draw on the experience of others (Metwally et al., 2024). Experience held by individual institutions is thereby converted into an artefact available to institutions that were not party to it, which is what distinguishes retention from continuity. The second is a feedback instrument: the second working group surveyed educational needs among members, and the results were reported as identifying where further material and professional development were required (Metwally et al., 2024). Materials are developed against the Agency's model academic curriculum, which exists in a revised edition (IAEA, 2021), and pass through Agency review before publication.

Three limits should be recorded. The account is a self-report by members of the network's leadership, and documents what exists more reliably than what it achieves. It is candid that much material remains undeveloped, and that although more than 65 per cent of survey respondents reported collaborating with other members, further collaboration is identified as an area for development, so the retention observed is partial. And the network operates within a state-centred and agency-supported frame. The case is closely related to the objectives of the resolution rather than conducted in its name.

The Tianjin Biosecurity Guidelines: translation across centres of knowledge production. The Tianjin Biosecurity Guidelines for Codes of Conduct for Scientists set out ten principles for the responsible conduct of the biological sciences. They were developed by Tianjin University's Center for Biosafety Research and Strategy, the Johns Hopkins Center for Health Security and the InterAcademy Partnership, which together facilitated two workshops with scientists from academia and industry across four continents; the work built on a working paper submitted to the Convention in 2016 and on efforts underway since 2015 (China & Pakistan, 2021, paras. 4–5). Roughly twenty international experts reviewed and revised earlier drafts, and the InterAcademy Partnership endorsed the Guidelines in July 2021 (InterAcademy Partnership, 2021). China and Pakistan submitted them to the Meeting of Experts on science and technology in September 2021, proposing that the Ninth Review Conference endorse them and that the intersessional process exchange information on their dissemination (China & Pakistan, 2021, para. 9). The case documents translation conducted between different centres of knowledge production, and diffusion through academy and treaty channels, sustained across more than one cycle. Its boundary is the one Table 1 predicts for diffusion: endorsement and submission are observable, adoption in laboratories is not, and a

voluntary text remains a voluntary text. The proposal to report on dissemination to a subsequent Review Conference is a design for feedback, not evidence that feedback has occurred. The case is closely related to the objectives of the resolution.

Read against the framework, the five cases distribute unevenly, and the unevenness is the result worth reporting. Partial connections are traceable, but no single case exhibits all four processes, and the chains break at identifiable points, in different places: at the boundary between an event and a network, at the join between co-present capacities, at the entrance to a retention channel, and at the end-point where diffusion meets adoption. The INSEN case provides direct evidence of diffusion and of institutional learning, and there a feedback mechanism can be pointed to rather than inferred; translation is more securely demonstrated in the Tianjin project coalition, and horizon sensing at the OPCW treaty interface, where assessment is an assigned function. Sensing is thus least documented in the network whose retention is strongest. In the two directly connected outreach cases the record establishes no academic-side retention and no subsequent uptake respectively, though neither absence should be generalised to Committee outreach as a whole.

Two consequences follow, and both narrow the claim this article makes. The first concerns the object: of the five cases, only one is unambiguously an academic knowledge network as defined in Section 3, the others being a treaty-facing advisory arrangement, a project coalition, a State-led outreach process and an event. The cases therefore identify complementary component capacities rather than a connected infrastructure, and the proposition defended here is correspondingly conditional.

The status of that result should be stated exactly, because it is easy to overstate. These cases were selected for variation in carrier type, region and boundary value, not to trace one configuration through all four processes, and they sit in different treaty regimes and technical fields. That they are unconnected to one another is therefore not evidence that the arrangements operating within any one field are unconnected. What the analysis establishes is narrower and, for a framework, more useful: demonstrating that component capacities co-exist does not demonstrate that an infrastructure exists; the latter claim requires traceable connection evidence of the kind specified in Section 4.1, and the present cases do not supply it. Whether it could be supplied is a question for research designed to follow one configuration, which this article has not attempted.

The second concerns coverage. African institutions participate substantially, forming one of the three largest regional groupings in the network examined above (Metwally et al., 2024), but no case is centred on an African institution, and participation in a network convened elsewhere is not a network with its centre of gravity in the region. Section 5 indicates where such carriers might be found; it does not close the gap.

5. Illustrative technology domains

Taking four broad areas of technological development in turn, this section asks what makes dual-use significance hard to see in each, and which communities are positioned to interpret it. Three of the four domains are supported by domain sources; the fourth is retained as a boundary question. The domains are treated as categories and the examples within them are illustrative.

Life sciences and biotechnology. Recognition in this domain often turns on context and intended use as much as on capability. Techniques that raise dual-use concern are frequently the same techniques that advance vaccine development or diagnostics, and concern may attach to the direction of an experiment rather than to the possession of a method. Early identification is therefore bound up with judgements about research design, to which researchers contribute knowledge that others do not hold. It does not follow that they hold it alone. The World Health Organization's global guidance framework treats the mitigation of biorisks and the governance of dual-use research as a shared responsibility, naming regulators, research institutions, educators, funding bodies, publishers, security actors and the private sector among those involved (World Health Organization [WHO], 2022). What the domain concentrates in researchers is a contribution to the judgement rather than the judgement entire,

and the interpreting communities extend accordingly. The Tianjin Guidelines belong here, and show a pattern that recurs below: translation accomplished among scientists, diffusion through academy channels, and what laboratories actually do left outside the record.

Artificial intelligence and data systems. Here the difficulty is different in kind. Dual-use significance may emerge from application rather than from design, and it may be demonstrated by the developers themselves rather than discovered by observers. A widely discussed illustration is the report by Urbina and colleagues that a generative model built for drug discovery, when its scoring function was inverted, produced large numbers of candidate toxic molecules within hours (Urbina et al., 2022). The example matters here for what it shows about recognition rather than for the capability itself: the finding came from the developers, was published in a machine-learning venue, and required biochemical knowledge to interpret. Sensing in this domain therefore appears to depend on technical and security knowledge that may not be institutionally co-located in any single community, which would strain the framework's first process. The Council has directed attention to rapid advances in science and technology and to information as an object of control (UNSC, 2016, paras. 7–8, 13), while naming no particular technology; this article follows it in that.

Advanced manufacturing and laboratory automation. What distinguishes this domain is that capability travels as design as well as as equipment. Physical controls continue to apply to machines, materials and feedstock, but a process encoded in a file is not a consignment, and to that extent the recognition problem is continuous with the intangible-transfer concern the resolutions have registered since 2011 (UNSC, 2011, preamble). Research on additive manufacturing has examined this combination directly, treating the control of machines, materials, software and design data as a single problem rather than as separable ones (Brockmann & Kelley, 2018). A further difficulty concerns the competence required to use those elements together, which diffuses through training, documentation and shared practice, that is, through exactly the channels the framework describes. This domain also shows the limits of the framework's academic focus, since much relevant expertise sits in industry rather than in universities, and industry is the community the resolution named first (UNSC, 2004, para. 8(d)).

Monitoring and detection technologies. Here the framework raises a question rather than answering one: whether the specialist communities that develop assurance technologies are also positioned to recognise developments that may erode those assurances. What would need to be seen early is not a hazardous capability but the weakening of an existing means of detection, which inverts the problem as posed in the other three domains. No case examined here answers that question, and no domain source is relied on, so the domain is retained as a boundary for future work rather than as evidence.

One point about the location of sensing should be kept separate from any particular domain. Where the assessment of scientific developments is an assigned institutional function, as it is for the scientific advisory mechanisms of treaty bodies, sensing is performed on a schedule and its outputs are documented. The OPCW's convergence assessment, discussed in Section 4, is an instance of that arrangement rather than of any single domain. Assigned sensing is regular and traceable, but it sits at a distance from the laboratories whose practice is at issue.

Read across the four domains, the examples suggest that the framework picks out different recognition problems rather than restating the same sequence in new vocabulary. Where dual-use significance turns on research design, sensing and translation may be closely coupled within research settings, although the resulting judgement remains distributed across other actors, and the difficulty lies further downstream, in whether anything is retained. Where significance emerges from combination across fields, sensing and translation come apart, and the knowledge needed to complete the judgement may not sit in one place. Where capability travels as information, the boundary of the academic sector itself becomes the constraint. These are different problems, and they suggest that the processes composing anticipatory awareness are unevenly distributed across domains.

A word is owed on the coverage gap identified in Section 4, which this section indicates but does not close. Federated biosafety associations are one place where relevant carriers might be sought: the International Federation of Biosafety Associations (IFBA) maintains a professional certification programme in biorisk

management structured in compliance with the international standard for bodies certifying persons, and national associations in Africa are among its members (IFBA, n.d.-a, n.d.-b). A certification scheme is a potential carrier of retention, since competence requirements may outlive those who meet them. But the available record does not establish when the federation was founded, how many associations belong to it, or whether its standards have been revised or used as a feedback mechanism, so the example documents no institutional learning; and African membership of an international federation is not a network centred on African institutions. That gap remains open.

If the difficulty of recognising dual-use significance takes a different form in each broad area of technological development, then the awareness-raising, education and knowledge-exchange activity that the resolutions associate with academia and industry (UNSC, 2016, preamble, paras. 29–30) cannot take a single form either; and since the materials, equipment and technology that States are required to control (UNSC, 2004, para. 3, footnote) are drawn from all of these areas, the unevenness described here bears on the environment in which the resolution's objective is pursued. Its relevance to that objective lies in helping the relevant communities recognise materials, equipment, technologies and knowledge of concern before States apply the awareness and control measures intended to reduce the pathways of access available to non-State actors; it does not itself constitute control or enforcement.

6. Limits and tensions

Five tensions are set out here. Each elaborates one or more of the failure modes entered in Table 1, though the relationship is analytical rather than one-to-one: a single tension may constrain several processes at once.

Openness and security. The processes described in Section 4 depend on knowledge circulating: sensing works because results are published and discussed, diffusion because materials can be shared. The same circulation can widen access to knowledge whose misuse would be of concern. The arrangements examined here do not eliminate that tension; they distribute its management across several sites of judgement. Poorly designed restrictions may narrow the expert scrutiny on which early recognition depends, while unrestricted circulation may widen the availability of sensitive knowledge. The World Health Organization's framework distributes the governance of dual-use research across many participants in the research life cycle rather than assigning it to one (WHO, 2022); the inference drawn here, which treats the preservation of multiple sites of judgement as a condition of effective sensing, is this article's rather than the framework's.

Fragmented expertise. The second tension is the one the framework was built to address. Translation requires a relationship or mechanism capable of integrating technical and non-proliferation judgements; it does not require that any one person hold both, which is why the arrangement matters more than the individual. Where no such relationship exists, the relevant judgements may remain distributed without being integrated. The OPCW case shows what an institutional answer looks like, a standing body whose task is to report scientific developments to a treaty community, and it also marks an evidentiary limit. Co-presence of scientific assessment and educational capacity cannot by itself establish a connection between them; documentation would have to trace an assessment into an educational product or practice first. The case therefore shows what demonstrating integration would require, not that integration has failed.

Unequal capacity. Participation is not capability. The network examined in Section 4 counts institutions in seventy-two States (Metwally et al., 2024), and a membership list of that reach tells us nothing about the resources, staffing or teaching infrastructure of any member. The same account reports that more than 65 per cent of survey respondents collaborated with other members, and identifies further collaboration as an area for development. Two consequences follow. Arrangements that appear global may in practice concentrate capacity in a subset of participants; and material developed under them may reflect the circumstances of the better-resourced, which limits its usefulness elsewhere. This article's own evidence base reproduces the problem, since cases were selected in part for documentary visibility, and the English-language record used here may favour institutions that publish detailed and readily accessible accounts of their work.

Over-securitisation. A framework that asks research communities to attend to dual-use significance carries a risk opposite to neglect. Definitions of concern that expand without limit may burden legitimate research, deter international collaboration, and be applied unevenly according to institutional affiliation or nationality. The resolution is explicit that preventing proliferation should not hamper international cooperation in materials, equipment and technology for peaceful purposes (UNSC, 2004, preamble), and the point is internal to the framework: an arrangement that flags too much loses the credibility on which its sensing depends, which is why over-flagging appears in Table 1 as a failure mode of the first process.

Voluntary participation and accountability. Participation in the academic and professional arrangements examined here is often voluntary, even where some interfaces operate within formally mandated international bodies. Voluntariness is the source of their flexibility and of three limits. Coverage is determined by who chooses to take part rather than by any principle of completeness. Outputs bind no one: a code of conduct endorsed by academies is a code of conduct endorsed by academies, and the Tianjin case shows how visible endorsement can be while adoption remains unmeasured. And non-mandated networks and project coalitions may depend heavily on temporary funding and individual commitment, which makes continuity and retention particularly fragile. These limits reinforce the importance of not confusing an epistemic supporting role with the legal obligations borne by States, which the resolutions place squarely on them and which academic expertise supports rather than shares (UNSC, 2016, para. 30).

Taken together, the five tensions suggest that the condition described here is not a stable achievement some communities possess and others lack, but a condition sustained against pressures that operate continuously and pull in different directions. That is a further reason for treating the proposition as conditional. How academic knowledge networks contribute to a distributed anticipatory infrastructure in any instance depends on how these tensions are managed and on whether their activities connect with those of other carriers; the framework specifies where to look rather than predicting what will be found. These tensions limit how academic arrangements can support awareness, knowledge exchange and early recognition relevant to the resolution's objective. They do not transfer or modify the obligations borne by States.

7. Implications for awareness and implementation-support

Three implications follow from the analysis. They are offered as observations about how the arrangements described might be understood, not as proposals addressed to any authority.

The first concerns what awareness-building consists of. The 2004 resolution connects working with and informing industry and the public to implementation; later resolutions explicitly extend civil-society outreach to academia (UNSC, 2004, para. 8(d); UNSC, 2016, preamble and para. 29; UNSC, 2022, para. 25). Across that development the raising of awareness is generally referred to as though it named a single activity. The framework suggests otherwise. Recognising a development, rendering it intelligible across a disciplinary boundary, circulating that understanding and retaining it are separable processes with different carriers and failure modes, so an activity addressing one leaves the others untouched. Training courses diffuse; they do not thereby sense, and a course whose materials and institutional follow-up are not retained leaves no documented carrier of durable diffusion. Awareness is better understood as a composite condition than as an output.

The second concerns translation. Within this small purposive set of cases, translation appears the least securely institutionalised of the four processes. Sensing has assigned homes in treaty-facing advisory bodies, and diffusion in universities and professional societies; retention was observable where a standing network codified its own experience. Translation, by contrast, appeared either as a function of an interface whose onward connection could not be traced, or as the achievement of a project assembled for the purpose and concluded. The process on which the others most depend may therefore be the one least likely to persist between occasions, since the persistence of project-shaped capacities cannot be assumed once the project ends.

The third concerns how the arrangements relate to one another. If the processes are distributed across carriers of different types, then no arrangement examined here is deficient for carrying only some, and the framework does not require any one arrangement to carry all four; that no single case documents a complete cycle describes an ecology rather than criticising its members. These arrangements may therefore be read more usefully as complementary than as alternatives, and the interesting questions concern the joins between them: whether an assessment produced in one setting reaches settings that could act on it, and whether anything is retained when it does. Those joins are where the documentary record was thinnest.

Two areas for further knowledge exchange follow. The connections between assessment, education and retention are easier to assert than to trace, and research following a particular scientific assessment through to its appearance, or non-appearance, in teaching materials and institutional practice would test the framework more severely than this article has been able to. Networks centred on institutions outside the best-documented regions are the second.

What these arrangements may affect is the epistemic condition in which the resolution's obligations are discharged: whether the materials, equipment and technology States are required to control are recognised as being of concern early enough for awareness, education and knowledge exchange to bear on them. That is a contribution upstream of implementation, and the resolutions locate academic expertise there, engaged in support of States and, where the Committee draws upon it, with State consent (UNSC, 2016, para. 30).

8. Conclusion

The question with which this article began was how academic knowledge networks contribute to anticipatory awareness of technological developments relevant to resolution 1540 (2004). The answer has three parts.

They contribute through identifiable processes rather than in general. Translation across disciplinary boundaries, diffusion through education and publication, and the retention of experience in forms outliving those who acquired it are distinct activities with distinct carriers, and the cases show them being performed. Where a standing network's working group contributed to codifying experience for institutions that had not shared in it, retention was observable as a mechanism rather than inferred from longevity.

They contribute partially. The component capacities are present, but no case exhibits the complete four-process configuration. Partial connections meeting the test set out in Section 4.1 are traceable, in the codification and feedback mechanisms of one network and in the passage of a jointly drafted text into academy channels, but not a chain running from sensing through to retention. Sensing was best evidenced where assessment is an assigned institutional function, at a distance from the laboratories concerned; retention was best evidenced in a network directed at educational provision rather than at technological signals. Academic knowledge networks are one component of that distribution, not an anticipatory infrastructure in themselves. Because the cases were chosen for variation rather than to follow a single configuration, this is a finding about what co-existence can establish, not a finding that the relevant arrangements are in fact unconnected.

And they contribute upstream of obligation. The duties created by resolution 1540 (2004) rest on States, and the arrangements examined here neither bear nor transfer them. What such arrangements may affect is whether the materials, equipment, technology and knowledge that States are required to control are recognised as being of concern early enough for awareness, education and knowledge exchange to bear on them, and thus, at a remove, on the pathways of access available to non-State actors that the resolution exists to close. Recognition is not control, and nothing described here performs the latter. The resolutions have themselves moved towards this ground, from a 2004 text that addressed industry and the public, through the recognition in 2011 that technology moves intangibly and that implementation is a long-term task, to the naming of academia in 2016 and the routinisation of review, feedback and outreach in 2022.

Two limits should close the account rather than qualify it from the margins. The evidence base favours arrangements that document themselves in accessible languages and formats, and no case examined here is centred on an African institution, despite substantial African participation in the network discussed in Section 4. And the condition described is not an achievement that some communities possess. It is a condition sustained against pressures, between openness and restriction, between specialisms and between unequal resources, that operate continuously and pull in different directions at different points.

Academic knowledge networks do not remove the uncertainties that attend emerging and dual-use technologies. Their significance lies in helping technical, academic and non-proliferation communities recognise, interpret and discuss those uncertainties at an early stage, and in retaining what is learned when they do. Whether that happens in any given instance is a question about arrangements, not about technologies, and it is answerable only by looking.

References

Barben, D., Fisher, E., Selin, C., & Guston, D. H. (2008). Anticipatory governance of nanotechnology: Foresight, engagement, and integration. In E. J. Hackett, O. Amsterdamska, M. Lynch, & J. Wajcman (Eds.), *The handbook of science and technology studies* (3rd ed., pp. 979–1000). MIT Press.

Brockmann, K., & Kelley, R. (2018). *The challenge of emerging technologies to non-proliferation efforts: Controlling additive manufacturing and intangible transfers of technology*. Stockholm International Peace Research Institute. https://www.sipri.org/sites/default/files/2018-04/sipri1804_3d_printing_brockmann.pdf

China & Pakistan. (2021). *The Tianjin biosecurity guidelines for codes of conduct for scientists* (BWC/MSP/2020/MX.2/WP.6 and Corr.1) [Working paper co-sponsored by Brazil, submitted to the Meeting of Experts on Review of Developments in the Field of Science and Technology, Geneva, 1–2 September 2021]. United Nations. https://docs.un.org/en/BWC/MSP/2020/MX.2/WP.6; https://docs.un.org/en/BWC/MSP/2020/MX.2/WP.6/Corr.1

Crossan, M. M., Lane, H. W., & White, R. E. (1999). An organizational learning framework: From intuition to institution. *Academy of Management Review, 24*(3), 522–537. https://doi.org/10.5465/amr.1999.2202135

1540 Committee. (n.d.). *Experiences shared, lessons learned and effective practices*. United Nations. Retrieved August 3, 2026, from https://www.un.org/en/sc/1540/national-implementation/experiences-shared-lessons-learned-and-effective-practices.shtml

1540 Committee. (2022). *Information note: Weapons of mass destruction, non-state actors and export control - Resolution 1540 (2004) of the Security Council and its national regional and global impact* [Seminar information note; event held 20–21 September 2022, Buenos Aires, Argentina]. United Nations. https://www.un.org/en/sc/1540/documents/Info-Note-CICTE-OAS-WMD-seminar-UBA-Argentina-2022-18.pdf

Guston, D. H. (2014). Understanding "anticipatory governance." *Social Studies of Science, 44*(2), 218–242. https://doi.org/10.1177/0306312713508669

Haas, P. M. (1992). Introduction: Epistemic communities and international policy coordination. *International Organization, 46*(1), 1–35. https://doi.org/10.1017/S0020818300001442

InterAcademy Partnership. (2021, July 8). *IAP endorses Tianjin Biosecurity Guidelines*. https://www.interacademies.org/news/iap-endorses-tianjin-biosecurity-guidelines

International Atomic Energy Agency. (2021). *Model academic curriculum in nuclear security* (IAEA Nuclear Security Series No. 12-T [Rev. 1]). https://www.iaea.org/publications/13608/model-academic-curriculum-in-nuclear-security

International Federation of Biosafety Associations. (n.d.-a). *Certification*. https://internationalbiosafety.org/certification/certification/

International Federation of Biosafety Associations. (n.d.-b). *Who we are*. Retrieved July 27, 2026, from https://internationalbiosafety.org/who-we-are/

Metwally, W. A., Gerlini, M., Ford, C., Fiel, J. C. B., & Goel, A. (2024). INSEN 2023 yearly updates. *International Journal of Nuclear Security, 9*(1), Article 3. https://doi.org/10.7290/ijns09735019

Organisation for the Prohibition of Chemical Weapons. (n.d.). *Advisory Board on Education and Outreach*. https://www.opcw.org/about/subsidiary-bodies/advisory-board-education-and-outreach

Organisation for the Prohibition of Chemical Weapons. (2014, June 30). *OPCW Scientific Advisory Board issues report on the convergence of chemistry and biology*. https://www.opcw.org/media-centre/news/2014/06/opcw-scientific-advisory-board-issues-report-convergence-chemistry-and

Permanent Mission of the Federal Republic of Germany to the United Nations. (2025, July 7). *Export controls: The importance of engaging academia* [Note verbale No. 279/2025 and annex]. United Nations Security Council Committee established pursuant to resolution 1540 (2004). https://www.un.org/en/sc/1540/documents/Germany1540Practices2025.pdf

Phelps, C., Heidl, R., & Wadhwa, A. (2012). Knowledge, networks, and knowledge networks: A review and research agenda. *Journal of Management, 38*(4), 1115–1166. https://doi.org/10.1177/0149206311432640

United Nations Security Council. (2004). *Resolution 1540 (2004)* (S/RES/1540 (2004)). https://docs.un.org/en/S/RES/1540(2004)

United Nations Security Council. (2011). *Resolution 1977 (2011)* (S/RES/1977 (2011)). https://docs.un.org/en/S/RES/1977(2011)

United Nations Security Council. (2016). *Resolution 2325 (2016)* (S/RES/2325 (2016)). https://docs.un.org/en/S/RES/2325(2016)

United Nations Security Council. (2022). *Resolution 2663 (2022)* (S/RES/2663 (2022)). https://docs.un.org/en/S/RES/2663(2022)

Urbina, F., Lentzos, F., Invernizzi, C., & Ekins, S. (2022). Dual use of artificial-intelligence-powered drug discovery. *Nature Machine Intelligence, 4*(3), 189–191. https://doi.org/10.1038/s42256-022-00465-9

World Health Organization. (2022). *Global guidance framework for the responsible use of the life sciences: Mitigating biorisks and governing dual-use research*. https://www.who.int/publications/i/item/9789240056107



1540 TechHorizon

Building a safer future
through science
technology
and global cooperation.

WARWICK
THE UNIVERSITY OF WARWICK

KING'S College LONDON

UCL

Cranfield University

UNODA
UNITED NATIONS OFFICE FOR DISARMAMENT AFFAIRS

Funded by the EU