frontiers in Synthetic Biology

TYPE Review
PUBLISHED 04 November 2026
DOI 10.3389/fsybi.2026.2005436

Check for updates

OPEN ACCESS

EDITED BY
David Bernard Levin,
University of Manitoba, Canada

REVIEWED BY
Ruimin Gao,
Public Health Agency of Canada (PHAC), Canada
Hazem Haddad,
Jordan University of Science and Technology, Jordan

*CORRESPONDENCE
Diego Cotella,
diego.cotella@uniupo.it

RECEIVED 21 September 2026
REVISED 25 October 2026
ACCEPTED 26 October 2026
PUBLISHED 04 November 2026

CITATION
Chawanda TV, Rimondini L and Cotella D (2026) Synthetic viruses: How far can we go?
Front. Synth. Biol. 4:2005436.
doi: 10.3389/fsybi.2026.2005436

COPYRIGHT
© 2026 Chawanda, Rimondini and Cotella. This is an open-access article distributed under the terms of the Creative Commons Attribution License (CC BY). The use, distribution or reproduction in other forums is permitted, provided the original author(s) and the copyright owner(s) are credited and that the original publication in this journal is cited, in accordance with accepted academic practice. No use, distribution or reproduction is permitted which does not comply with these terms.

Synthetic viruses: How far can we go?

Tonderai Vitalis Chawanda, Lia Rimondini and Diego Cotella*

Department of Health Sciences, Università Del Piemonte Orientale, Novara, Italy

Advances in molecular biology, DNA synthesis, and computational design have made it possible to build viruses from written genetic sequence alone, without a natural template. Since the first chemical synthesis of an infectious poliovirus genome in 2002, synthetic virology has produced steadily larger and more complex genomes, reconstructed extinct pandemic viruses for research, and assembled a live orthopoxvirus from mail-order DNA fragments. In parallel, the same underlying toolkit, virus reverse-genetics systems, genome and gene synthesis, and increasingly artificial intelligence-guided sequence and protein design, has produced major public health benefits, including messenger RNA and self-amplifying RNA vaccines, oncolytic virotherapies and engineered gene therapy vectors. This review surveys the history, enabling technologies, and current capabilities of synthetic virology; describes its principal applications in vaccinology, oncology, and gene therapy; and examines the dual-use biosecurity risks this capability creates, including gain-of-function research, the declining cost and increasing accessibility of gene synthesis, and the emerging challenge of AI-designed sequences evading biosecurity screening. It then reviews the governance landscape, including dual-use research of concern oversight, nucleic acid synthesis screening, and international coordination efforts, before considering how far the field can and should go, and what adaptive governance would be required to keep pace with it.

KEYWORDS
artificial intelligence in biology, biosecurity, *de novo* virus synthesis, DNA synthesis screening, dual-use research of concern, gain-of-function research, mRNA vaccines, oncolytic virotherapy

1 Introduction

In July 2002, a research team at Stony Brook University reported something that had never been done before: they had built an infectious virus from chemicals alone. Working only from the published genetic sequence of poliovirus, the team chemically synthesized overlapping DNA fragments, assembled them into a complete complementary DNA copy of the viral genome, transcribed that DNA into RNA, and showed that the resulting RNA was infectious in a cell-free extract, producing virus particles indistinguishable in essential respects from naturally occurring poliovirus (Cello et al., 2002). No natural template, no infected cell, no viral stock was needed at any point; the starting material was a written sequence and a mail-order supply of oligonucleotides. One of the authors later described the result in stark terms, noting that eradicating a virus from nature no longer guarantees that it cannot return, because "somebody can always resynthesize it" (Wimmer, 2006).

That single experiment opened what is now called synthetic virology: the design, construction, and modification of viral genomes using tools borrowed from chemistry, molecular biology, and, increasingly, computer science and artificial intelligence. In the 2 decades since, the field has moved from proof-of-concept curiosities to core infrastructure of modern biomedicine. The same reverse-genetics and genome-synthesis toolkit that let

Cello, Paul, and Wimmer build poliovirus from scratch now underlies the messenger RNA (mRNA) vaccine platforms that were deployed within weeks of the SARS-CoV-2 genome sequence becoming public, the oncolytic viruses that are approved cancer therapies, and the viral vectors used in gene therapy. Yet the very features that make synthetic virology powerful for medicine, namely, that a virus can be built from a sequence rather than isolated from nature, that costs of DNA synthesis keep falling, and that computational design tools are becoming ever more capable, also make it a recognized biosecurity concern. The 2018 report that a poxvirus had been chemically synthesized and assembled into a live, replicating orthopoxvirus made that tension explicit and reignited a debate that has never fully settled: how far should this capability be allowed to go, and who decides (Noyce and Evans, 2018; Noyce et al., 2018)?

This review asks that question in three parts. First, it traces the history and technical foundations of synthetic virology, from the earliest chemical synthesis of a viral genome to the largest genome assembled to date, in order to establish what "building a virus" actually involves and how the achievable scale and complexity have grown. Second, it surveys the beneficial applications this capability has enabled, in vaccinology, oncology, and gene therapy, to show what has already been gained. Third, it examines the risks: dual-use research of concern (DURC), gain-of-function (GOF) studies on pandemic-potential pathogens, the declining cost and increasing accessibility of DNA synthesis, and the emerging challenge that artificial intelligence poses to the biosecurity screening systems meant to catch dangerous orders before they are filled. It then reviews the governance architecture that has grown up around this dual-use problem and closes by considering, in light of all of the above, how far synthetic virology can realistically go, and how far it should.

2 A brief history of synthetic virology

Although the 2002 poliovirus synthesis is usually cited as the field's founding experiment, it built on more than 2 decades of reverse-genetics work that let researchers derive infectious virus from cloned DNA rather than from a biological sample, beginning with the molecular cloning and full sequencing of the poliovirus genome in 1981 (Racaniello and Baltimore, 1981). What made the 2002 study distinctive was that the DNA itself, not just the recovery method, was built from chemical building blocks rather than copied from an existing template, closing the loop from written sequence to living pathogen. The scale was modest by later standards: the poliovirus genome is a single-stranded RNA molecule of about 7,500 nucleotides, among the smallest of clinically important viral genomes (Cello et al., 2002).

The following year, a team at the J. Craig Venter Institute (JCVI) demonstrated a faster and more general assembly method by synthesizing the roughly 5,400 base pair genome of the bacteriophage phiX174 from a pool of chemically synthesized oligonucleotides, assembling the complete genome in about 2 weeks and showing that speed, not just feasibility, would define the next phase of the field (Smith et al., 2003). In 2005, synthetic virology was turned toward a very different purpose: understanding a historical pandemic. Researchers used sequence recovered from archived tissue samples to reconstruct the complete genome of the 1918 influenza virus, responsible for the deadliest pandemic of the 20th century, and generated live virus bearing all eight of its original gene segments in order to study the molecular basis of its exceptional virulence (Tumpey et al., 2005). Reconstructing an extinct pandemic virus for research, inside controlled biocontainment, illustrated both the scientific value of synthetic virology and the fact that "extinction" of a pathogen no longer means it is permanently unavailable, whether for legitimate research or otherwise.

The following years extended synthesis from viral genomes to whole cellular genomes. In 2008, the same Venter Institute team chemically synthesized, assembled, and cloned the complete roughly 583-kilobase-pair genome of the bacterium *Mycoplasma genitalium*, at the time the largest chemically defined structure ever built, establishing hierarchical assembly strategies that synthetic virology would later borrow for large DNA viruses (Gibson et al., 2008). In 2016, the same group combined whole-genome design with complete chemical synthesis to build JCVI-syn3.0, a minimal bacterial genome of 531 kilobase pairs and 473 genes, smaller than that of any known autonomously replicating cell (Hutchison et al., 2016). Although both are bacteria rather than viruses, this design-build-test-learn (DBTL) cycle, in which an initial computational design is synthesized, tested for viability, and iteratively refined, established a general workflow for whole-genome engineering that has since informed viral genome design as well, and demonstrated that genomes an order of magnitude larger than early synthetic viruses could be designed and built essentially from scratch.

The event that most directly reopened public debate about the limits of synthetic virology came in January 2018, when a Canadian research team published the complete synthesis of an infectious horsepox virus, assembled from ten large chemically synthesized DNA fragments recombined inside cells infected with a helper virus (Noyce et al., 2018). At roughly 212 kilobase pairs, the resulting synthetic chimeric horsepox virus (scHPXV) was, at the time, the largest virus genome ever assembled through synthetic biology, and it belonged to the orthopoxvirus genus that also contains variola virus, the causative agent of smallpox (Coyne, 2018; Noyce et al., 2018). The researchers' stated purpose was benign: horsepox, believed extinct in nature since the 1980s and unavailable as a physical specimen, was reconstructed as a candidate platform for a safer smallpox vaccine, motivated by evidence that Edward Jenner's original cowpox-derived vaccine may in fact have been horsepox-derived (Noyce et al., 2018). But because any method capable of assembling one orthopoxvirus from synthetic DNA fragments is, in principle, capable of

Abbreviations: AAV, adeno-associated virus; AI, artificial intelligence; BWC, Biological Weapons Convention; cDNA, complementary DNA; DBTL, design–build–test–learn; DURC, dual-use research of concern; GM-CSF, granulocyte-macrophage colony-stimulating factor; GOF, gain-of-function; HPTA, Human Pathogens and Toxins Act; HSV-1, herpes simplex virus type 1; IBBIS, International Biosecurity and Biosafety Initiative for Science; IGSC, International Gene Synthesis Consortium; JCVI, J. Craig Venter Institute; mRNA, messenger RNA; NSABB, National Science Advisory Board for Biosecurity; P3CO, Potential Pandemic Pathogen Care and Oversight; saRNA, self-amplifying RNA; scHPXV, synthetic chimeric horsepox virus; T-VEC, talimogene laherparepvec; WHO, World Health Organization.

TABLE 1 Selected milestones in synthetic virology and closely related whole-genome synthesis, in chronological order.

Year	Milestone	Ref.
1981	First molecular cloning and full sequencing of a viral (poliovirus) genome, enabling reverse genetics	(Racaniello and Baltimore, 1981)
2002	First infectious virus (poliovirus, ~7.5 kb) built entirely from chemically synthesized DNA.	(Cello et al., 2002)
2003	Bacteriophage ϕX174 genome (~5.4 kb) assembled from synthetic oligonucleotides in ~2 weeks	(Smith et al., 2003)
2005	Complete 8-segment genome of the 1918 pandemic influenza virus reconstructed and recovered as live virus	(Tumpey et al., 2005)
2008	Complete *Mycoplasma* genitalium bacterial genome (~583 kb) chemically synthesized, assembled, and cloned	(Gibson et al., 2008)
2009	Isothermal, single-reaction (Gibson) assembly method enables seamless joining of large DNA fragments	(Gibson et al., 2009)
2016	JCVI-syn3.0 minimal bacterial genome (531 kb, 473 genes) built via iterative design-build-test cycle	(Hutchison et al., 2016)
2018	Largest virus genome synthesized to date: ~212 kb synthetic chimeric horsepox virus assembled from 10 DNA fragments	(Noyce et al., 2018)
2025	AI protein-design tools shown to generate sequence variants that evade biosecurity screening; patches deployed	(Wittmann et al., 2025)
2026	First generative design of a complete, functional bacteriophage genome (ϕX174-derived) using genome language models; 16 viable AI-designed phages produced	(King et al., 2026)

assembling another, the study was immediately recognized as a case of dual-use research of concern, research undertaken for legitimate reasons whose methods and results could also be used for harm (DiEuliis et al., 2017; Koblentz, 2017; Noyce and Evans, 2018). The horsepox synthesis, together with the reconstruction of the 1918 influenza virus and the earlier poliovirus and phiX174 syntheses, forms the core sequence of milestones against which the current reach of synthetic virology is usually measured; Table 1 summarizes this trajectory, and Section 4 discusses what it implies about the field's present ceiling.

3 Enabling technologies for building viruses

Building a virus from sequence rather than isolating it from nature depends on a handful of underlying technologies, each of which has matured considerably since 2002. This section describes them at a conceptual level, sufficient to understand what "synthesizing a virus" entails and why the field's reach has expanded, without providing procedural detail that would be needed to reproduce any specific pathogen.

3.1 Reverse genetics

Reverse genetics refers to the general strategy of starting from a defined, cloned copy of a viral genome, rather than from an infected host or cell culture, and using it to recover infectious virus. The approach traces to the early 1980s, when researchers first cloned poliovirus complementary DNA and determined the full nucleotide sequence of the viral genome, showing that a plasmid-borne DNA copy could substitute for the natural RNA genome as the starting point for recovering live virus (Racaniello and Baltimore, 1981). Reverse-genetics systems were subsequently developed for other RNA and DNA virus families through the 1980s and 1990s, originally using DNA copied enzymatically from natural viral RNA or DNA. Synthetic virology extends this logic one step further by replacing the natural template with a chemically synthesized one, so that the starting point is a digital sequence file rather than a physical biological sample (Cello et al., 2002). The technical distinction matters: a reverse-genetics system built from a natural template can only recover viruses that someone has already isolated, whereas one built from synthetic DNA can in principle recover any virus whose sequence has ever been determined and published, including viruses that are extinct, unavailable, or that exist only as a designed variant.

3.2 Gene and genome synthesis and assembly

Modern gene synthesis builds long DNA molecules from short, chemically synthesized oligonucleotides, which are joined into progressively larger fragments and finally into a complete gene or genome using enzymatic assembly methods. A key enabling advance was the development of isothermal, single-reaction assembly methods, commonly known as Gibson assembly, in which a 5′exonuclease, a DNA polymerase, and a DNA ligase act together to join multiple overlapping DNA fragments seamlessly and without the sequence restrictions imposed by traditional restriction-enzyme cloning (Gibson et al., 2009) (Figure 1). This and related methods allowed the scale of what could be assembled in a single project to grow from single genes to genomes exceeding 200 kilobase pairs and were used directly in both the 2008 synthesis of the *M. genitalium* genome and the 2018 horsepox synthesis (Gibson et al., 2008; Noyce et al., 2018). Whole-genome design projects such as JCVI-syn3.0 additionally use a design-build-test-learn cycle: a computer-generated design is synthesized, tested for function, and revised based on the result, allowing genomes to be engineered rather than merely copied (Hutchison et al., 2016). As the cost of DNA synthesis has fallen and turnaround times have shortened, gene and genome synthesis has moved from a specialized academic capability to a

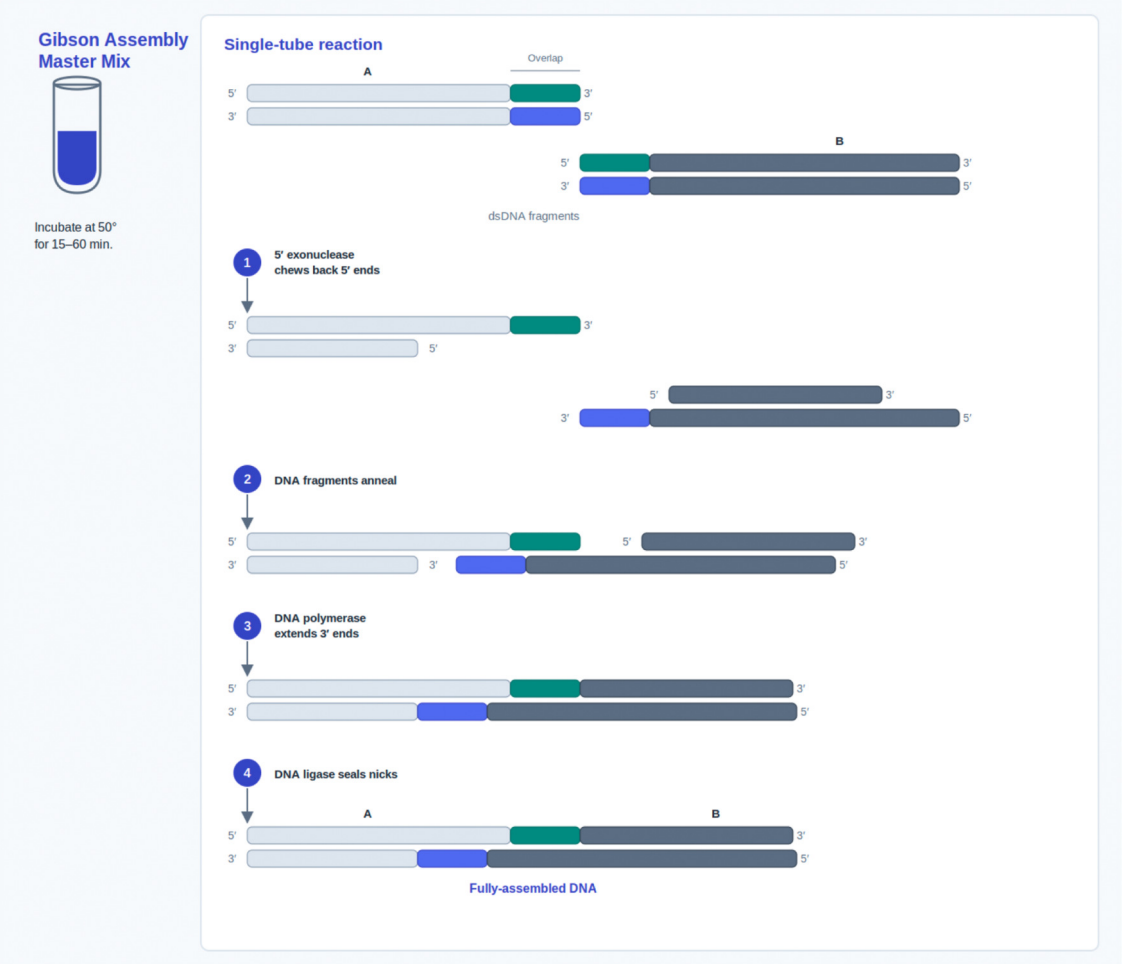


FIGURE 1
Mechanism of Gibson assembly. Two DNA fragments sharing 20–40 base pairs of identical sequence at their ends (1) are treated with a 5′exonuclease, which resects the 5′ends and exposes complementary single-stranded 3′overhangs (2). The overhangs anneal, joining the fragments while small nicks remain in the backbone (3). A DNA polymerase fills any remaining single-stranded gap and a DNA ligase seals the nicks, yielding one continuous double-stranded molecule (4). All four activities proceed in a single isothermal reaction (Gibson et al., 2009).

commercial service offered by dozens of providers worldwide, a trend discussed further in Section 6.3.

3.3 Directed evolution and rational design

Not every improvement to a viral or viral-derived protein comes from designing a sequence from first principles. An alternative strategy, directed evolution, mimics natural selection in the laboratory: a gene is randomly mutated and recombined, the resulting variants are expressed, and those with improved or altered function are selected and subjected to further rounds of mutation and selection. The approach, recognized with the 2018 Nobel Prize in Chemistry, was originally developed to engineer enzymes for industrial and pharmaceutical use but has since been applied broadly across protein engineering, including to viral capsid and vector proteins, where it can identify functional variants that would be difficult to predict from structure alone (Nobel Prize in Chemistry, 2018). Directed evolution and rational, computational design (Section 3.4) are complementary rather than competing strategies: modern viral vector engineering, discussed in Section 5.3, often combines an initial rationally designed scaffold with subsequent rounds of directed evolution to optimize properties such as tissue tropism or immune evasion that remain difficult to predict computationally.

3.4 Computational and AI-Assisted design

The most recent addition to the synthetic virology toolkit is computational design, and in particular the use of deep-learning models to design proteins and nucleic acid sequences with specified structures or functions. Diffusion-based generative models such as RFdiffusion can generate novel protein backbones and sequences for therapeutic and industrial applications, including binder design, enzyme scaffolding, and *de novo* protein architectures with no natural counterpart (Watson et al., 2023). Applied to viral proteins, comparable tools can in principle generate sequence variants that retain a protein's function or structural fold while differing substantially from any natural or previously characterized sequence, a capability with direct implications for both vaccine design and biosecurity screening, discussed in Section 6.4 (Watson

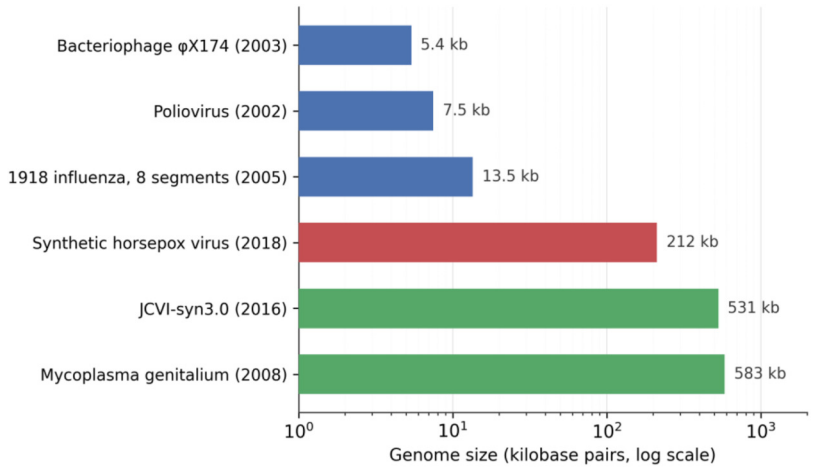


FIGURE 2
Genome size of selected synthetic-biology milestones, plotted on a logarithmic scale. The 2018 synthetic chimeric horsepox virus (red) remains the largest virus genome assembled to date, roughly thirty times the size of the 2002 synthetic poliovirus genome; the two whole-genome bacterial projects (green), which used closely related assembly methods, show that several-times-larger genomes are already within reach of the same toolkit (Cello et al., 2002; Smith et al., 2003; Tumpey et al., 2005; Gibson et al., 2008; Hutchison et al., 2016; Noyce et al., 2018).

et al., 2023; Baker and Church, 2024). A 2026 study shows how far this approach has already advanced beyond individual proteins to entire viral genomes. Using genome language models, Evo one and Evo 2, trained on natural sequence diversity, researchers generated complete, functional bacteriophage genomes rather than isolated genes or proteins, using the roughly 5.4 kilobase genome of bacteriophage phiX174, the same phage synthesized chemically in 2003 (Section 2), as a design template (King et al., 2026). Sixteen of the generated genomes proved viable when introduced into host *Escherichia coli*, producing phage particles with substantial sequence novelty relative to any known natural phage; cryo-electron microscopy of one generated phage revealed a DNA-packaging protein evolutionarily distant from anything found in phiX174 or its known relatives, indicating that the model had accessed regions of functional sequence space not represented in its training data (King et al., 2026). The result is, to date, the clearest demonstration that generative sequence models can design an entire functional virus genome rather than a single gene or protein, extending the computational-design toolkit described in this section from the protein scale to the genome scale and materially sharpening the biosecurity questions discussed in Section 6.4 (Inglesby and Hanke, 2026; King et al., 2026).

4 How far has the field already gone? Current capabilities

Taken together, the milestones described in Section 2 chart a fairly clear trajectory of increasing genome size, decreasing turnaround time, and increasing design sophistication, summarized in Table 1 and Figure 2.

Two conclusions follow. First, no single technical barrier separates small RNA viruses from much larger DNA viruses or even cellular genomes; the horsepox synthesis showed that a poxvirus genome roughly thirty times the size of the poliovirus genome could be assembled from commercially synthesized fragments using published methods, and the 2008 and 2016 *Mycoplasma* genome projects showed that similar methods scale to genomes several times larger still (Gibson et al., 2008; Hutchison et al., 2016; Noyce et al., 2018). As the authors of the horsepox study later wrote, reflecting on the implications of their own work, the result indicates that "no viral pathogen is likely beyond the reach of synthetic biology" (Noyce and Evans, 2018). Second, the rate-limiting step has shifted over time: what once required specialized academic infrastructure now depends primarily on commercially available gene synthesis, meaning that the practical limit on what can be built is increasingly a function of cost, screening, and oversight rather than raw technical capability. That shift is the central concern of Sections 6 and 7.

It is worth being precise about what these milestones do, and do not, establish. None of them created a novel pathogen with capabilities exceeding those of known viruses; each recreated, reconstructed, or reorganized genetic material corresponding to an existing or historical virus. The scientific and biosecurity significance lies not in any single result but in the method: each demonstrated that a virus, or virus-sized genome, of a given size and complexity could be built from synthetic DNA and a written sequence, establishing a technical benchmark that lowers the barrier for the next such effort, whether beneficial or otherwise.

5 Applications: Turning synthesis into medicine

The capabilities described above are not solely, or even primarily, biosecurity liabilities. The same core toolkit, reverse genetics, gene and genome synthesis, directed evolution, and computational sequence design, underlies several of the most important recent advances in vaccinology and virus-based

TABLE 2 Selected synthetic-virology applications in vaccinology, oncology, and gene and phage therapy.

Application	Platform/mechanism	Example product	Approval/status	Ref.
mRNA vaccine	Lipid-nanoparticle-delivered mRNA encoding a pathogen antigen	BNT162b2 (Pfizer–BioNTech); mRNA-1273 (Moderna)	FDA/EMA approved, 2020–2021	(Polack et al., 2020; Baden et al., 2021)
Self-amplifying RNA vaccine	mRNA + alphavirus-derived replicase; self-amplifies intracellularly	ARCT-154 (Arcturus/CSL)	Approved, Japan, Nov. 2023	(Hồ et al., 2024)
Synthetic chimeric vaccine virus	*De novo* assembly of an attenuated orthopoxvirus from synthetic DNA	Synthetic chimeric horsepox virus (scHPXV)	Preclinical candidate	(Noyce et al., 2018)
Oncolytic virotherapy	HSV-1 engineered for tumor-selective replication + GM-CSF expression	Talimogene laherparepvec (T-VEC, Imlygic)	FDA approved 2015; EU, Australia, Israel	(Amgen, 2015; Andtbacka et al., 2015)
Oncolytic virotherapy	E1B-deleted oncolytic adenovirus	Oncorine	Approved, China	(Kaufman, 2022)
Oncolytic virotherapy	Unmodified oncolytic picornavirus	Rigvir	Approved, parts of Eastern Europe	(Kaufman, 2022)
Oncolytic virotherapy	Triple-mutated oncolytic HSV	Teserpaturev (Delytact)	Approved, Japan, 2021	(Kaufman, 2022)
Gene therapy vector	Replication-deficient AAV/lentiviral vectors	Vectors for inherited retinal disease, spinal muscular atrophy, hemophilia	Multiple approved products	(High and Roncarolo, 2019)
Phage therapy	Natural or AI-designed bacteriophage cocktails targeting resistant bacteria	AI-designed φX174-derived phage cocktail (proof of concept)	Case reports/early trials; research stage	(Kortright et al., 2019; King et al., 2026)

therapeutics; Table 2 summarizes the platforms, representative products, and regulatory status discussed in this section.

5.1 Vaccines: mRNA and self-amplifying RNA platforms

Messenger RNA vaccines are themselves synthetic biology products: a computer-designed sequence encoding a pathogen antigen is chemically synthesized and packaged for delivery, without any need to grow, attenuate, or inactivate the pathogen itself. The platform's clinical viability rests on a discovery made 2 decades earlier than the COVID-19 pandemic: that substituting modified nucleosides into synthetic mRNA suppresses the innate immune activation that had previously made therapeutic mRNA impractical, work recognized with the 2023 Nobel Prize in Physiology or Medicine (Karikó et al., 2005; Nobel Prize in Physiology or Nobel Prize in Physiology or Medicine, 2023). Built on that foundation, the Pfizer-BioNTech and Moderna mRNA vaccines each demonstrated approximately 95 percent efficacy against symptomatic COVID-19 in large randomized phase III trials completed within about a year of the virus's genome sequence becoming public (Polack et al., 2020; Baden et al., 2021), a timeline that would have been essentially unimaginable for a conventional, cell-culture-grown vaccine platform.

Self-amplifying RNA (saRNA) vaccines extend this approach by encoding, alongside the antigen, a viral replicase borrowed from an alphavirus genome, allowing the RNA to copy itself once inside a cell and produce substantially more antigen from a much smaller dose than conventional mRNA vaccines (Bloom et al., 2021; Casmil et al., 2025). Because the replicase genes retain none of the structural genes needed to build new virus particles, saRNA vaccines are able to amplify locally without producing infectious virus (Casmil et al., 2025). In November 2023, Japan's regulatory authority approved ARCT-154, developed by Arcturus Therapeutics and CSL, as the first self-amplifying mRNA vaccine authorized for use against COVID-19, with pooled clinical trial data showing immune responses comparable to conventional mRNA vaccines at a fraction of the dose (Hồ et al., 2024). Self-amplifying and circular RNA platforms are now being explored for a widening range of infectious diseases, including influenza and emerging zoonotic threats, and for applications beyond infectious disease such as cancer immunotherapy (Bloom et al., 2021; Casmil et al., 2025).

Synthetic virology also enables rational vaccine design of a more direct kind: building an entire attenuated or chimeric virus, rather than just an RNA-encoded antigen, as a vaccine platform. The horsepox synthesis discussed above was itself undertaken for this purpose. The resulting synthetic chimeric horsepox virus produced smaller plaques, less extracellular virus, and reduced virulence in mice compared with modern vaccinia-based smallpox vaccines, while still conferring protection against a lethal vaccinia challenge, supporting its further development as a next-generation smallpox vaccine (Noyce et al., 2018). The episode illustrates a recurring pattern in this field: the same construction method that raises a biosecurity concern is often motivated by, and capable of producing, a genuine public health benefit.

A related genome-driven methodology, reverse vaccinology, extended this logic from vaccine platforms to antigen discovery itself. Rather than growing a pathogen and testing its proteins individually for immunogenicity, reverse vaccinology starts from a complete genome sequence and uses computational analysis to predict which surface-exposed or secreted proteins are most likely to be accessible to the immune system, producing a shortlist of candidate antigens for experimental validation before any of the organism is cultured at scale. Rino Rappuoli pioneered the approach and coined the term in the late 1990s, applying it to *Neisseria*

meningitidis serogroup B, a bacterium that had resisted decades of conventional vaccine development because its principal surface polysaccharide is a poor immunogen and structurally similar to a human neural antigen (Pizza et al., 2000; Rappuoli, 2000). The project relied on a complete genome sequence generated by The Institute for Genomic Research, co-founded by J. Craig Venter, whose sequencing infrastructure exemplified the same genome-centric approach that, within a few years, would also underpin the phiX174 and *M. genitalium* synthesis projects described in Section 2 (Pizza et al., 2000; Tettelin et al., 2000). Computational screening of the genome yielded several hundred candidate antigens, a subset of which were validated experimentally and ultimately combined into Bexsero, a licensed multicomponent meningococcal serogroup B vaccine (Pizza et al., 2000). Reverse vaccinology illustrates that the genome-centric toolkit described throughout this review has reshaped not only how vaccine platforms and viral vectors are built, but how the antigens within them are identified in the first place. This same genome-to-antigen logic proved decisive during the COVID-19 pandemic. Because reverse vaccinology and structure-based antigen design both start from a genetic sequence rather than from a cultured pathogen, prior structural work on other betacoronavirus spike proteins, developed through years of MERS-CoV research, allowed the prefusion-stabilizing mutations needed for an effective SARS-CoV-2 antigen to be identified computationally within days of the viral genome sequence becoming public in January 2020, without isolating or culturing the virus itself (Corbett et al., 2020). That antigen design was handed directly to Moderna for mRNA vaccine manufacturing, and the resulting candidate, mRNA-1273, entered a first-in-human clinical trial roughly 2 months after the sequence was released, a timeline that its developers explicitly credited to this "prototype pathogen preparedness" strategy, in effect the extension of reverse vaccinology's genome-first logic from bacterial antigen discovery to rapid viral antigen design (Graham, 2020).

5.2 Oncolytic virotherapy

Oncolytic viruses are engineered or naturally occurring viruses that preferentially infect and destroy cancer cells while sparing healthy tissue, often while also stimulating an anti-tumor immune response. Talimogene laherparepvec (T-VEC, marketed as Imlygic), a herpes simplex virus type 1 engineered to replicate selectively in tumor cells and to express granulocyte-macrophage colony-stimulating factor to boost anti-tumor immunity, became the first oncolytic virus therapy approved by the United States Food and Drug Administration in October 2015, for the local treatment of advanced melanoma (Amgen, 2015; Andtbacka et al., 2015). Approval followed a phase III trial in which T-VEC produced a significantly higher durable response rate than a comparator cytokine therapy (Andtbacka et al., 2015). T-VEC has since been approved in Europe, Australia, and Israel, and related oncolytic viruses have reached the market elsewhere: an E1B-deleted oncolytic adenovirus (Oncorine) is approved in China for head and neck cancers, an unmodified picornavirus (Rigvir) is approved for melanoma in parts of Eastern Europe, and a triple-mutated oncolytic herpesvirus (teserpaturev, Delytact) was approved in Japan in 2021 for malignant glioma (Kaufman, 2022). Oncolytic virotherapy illustrates a distinct application of the synthetic virology toolkit: rather than reconstructing a pathogen exactly as it occurs in nature, it deliberately re-engineers a virus's tropism and gene expression to redirect its natural ability to infect and kill cells toward a therapeutic target.

5.3 Gene therapy and engineered vectors

Further application lies in the use of engineered, typically replication-deficient, viruses as delivery vehicles for gene therapy, in which a virus's natural ability to enter cells and deliver genetic material is repurposed to deliver a therapeutic gene rather than to replicate itself. Adeno-associated virus and lentiviral vectors, engineered so that they cannot reproduce independently, are now used clinically to deliver corrective genes for inherited disorders such as inherited retinal disease, spinal muscular atrophy, and hemophilia, and are being investigated for a wide range of additional indications (High and Roncarolo, 2019). The design principles are continuous with those described in Section 3: a viral genome is redesigned at the sequence level, synthesized, and tested, often using rounds of directed evolution to refine capsid properties such as tissue tropism, with the specific aim of removing the pathogenic and replicative machinery while retaining only the components needed for safe and effective delivery (High and Roncarolo, 2019).

5.4 Phage therapy

A related, bacteria-targeting application has seen renewed clinical interest as antibiotic resistance has spread: bacteriophage (phage) therapy uses viruses that infect and lyse specific bacteria as an alternative or adjunct to conventional antibiotics. Long overshadowed by the rise of chemical antibiotics after their discovery in the early 20th century, phage therapy has been revived through case reports and early clinical trials demonstrating utility against otherwise intractable multidrug-resistant infections, and through the same genome-engineering toolkit described in Section 3, which allows phages to be selected, evolved, or redesigned to target resistant strains and to be combined rationally with antibiotics (Kortright et al., 2019). Phage therapy underscores that synthetic virology's reach extends beyond viruses of humans and animals to the far larger population of viruses that infect bacteria, with direct relevance to the separate global health challenge of antimicrobial resistance. The generative genome-design approach described in Section 3.4 has already been tested directly against this problem: a cocktail of AI-designed, phiX174-derived phages was shown to rapidly overcome *E. coli* strains that had evolved resistance to the natural parent phage, a result its authors present as an early proof of concept for AI-generated phage therapeutics capable of keeping pace with fast-evolving, drug-resistant bacteria (King et al., 2026). Because the approach designs at the level of a complete genome rather than a single resistance-conferring gene, it points toward phage cocktails that could in principle be refreshed computationally as bacteria evolve resistance to earlier phage generations, rather than relying solely on isolating new phages from the environment.

6 The other edge: Risks and dual-use concerns

The same capabilities described in Sections 3 through 5 are, by their nature, dual-use: the tools that allow a candidate vaccine virus to be designed and built are the same tools that would be needed to reconstruct a dangerous pathogen. This section examines four aspects of that risk: the formal concept of dual-use research of concern and the horsepox case that crystallized it; the related but distinct controversy over gain-of-function research; the economic and logistical trend of falling synthesis costs and rising accessibility; and the newest complication, the use of artificial intelligence to design sequences that may evade the screening systems built to catch them.

6.1 Dual-use research of concern and the horsepox case

Dual-use research of concern (DURC) refers to legitimate, often clearly beneficial, life-sciences research whose knowledge, products, or methods could be readily misused to pose a significant biological threat. The 2018 horsepox synthesis is the case most often cited to illustrate the concept in synthetic virology, precisely because the researchers' goal, a safer smallpox vaccine platform, was legitimate and well-motivated, while the method they published, *de novo* assembly of a large orthopoxvirus genome from commercially synthesized DNA fragments, is generalizable to variola virus itself (DiEuliis et al., 2017; Koblentz, 2017; Noyce and Evans, 2018; Noyce et al., 2018). Critics argued that the study's methodology provided a detailed roadmap that meaningfully lowered the technical and logistical barriers to reconstructing smallpox, a disease eradicated from nature in 1980 whose only remaining declared stocks are held under strict security in two national laboratories, and that the risk of publication outweighed the benefit of an improved vaccine candidate for a disease that, absent reintroduction, no longer circulates (DiEuliis et al., 2017; Koblentz, 2017). The study's authors, and the editor who oversaw its review, countered that the underlying vulnerability, namely, that gene synthesis technology could in principle be used to reconstruct an eradicated orthopoxvirus, already existed independently of their publication, that responsible disclosure allows the public health and biosecurity communities to prepare defenses, and that historical experience with genetic engineering suggests that technologies initially viewed as threats have often gone on to substantially improve human welfare (Coyne, 2018; Noyce and Evans, 2018). Both positions are represented in the biosecurity literature, and the disagreement is as much about how to weigh a diffuse, long-term proliferation risk against a concrete, near-term scientific benefit as it is about the specific facts of the case.

The dual-use character of this field is best understood through documented cases spanning three distinct hazard types. The first is the unintended enhancement of virulence: researchers developing a mouse-contraceptive vaccine inserted the interleukin-4 gene into ectromelia (mousepox) virus and inadvertently produced a variant that suppressed the antiviral cell-mediated immune response and proved lethal even in vaccinated and genetically resistant animals, demonstrating that a routine, benignly motivated genetic modification can unexpectedly increase lethality (Jackson et al., 2001). The second is the sequence-level information hazard: when a novel botulinum neurotoxin for which no antitoxin existed was characterized in 2013, the responsible authors and journals, in consultation with public-health authorities, withheld the toxin's gene and protein sequences from publication until countermeasures could be developed, the first instance in which a peer-reviewed article deliberately omitted a biological sequence for security reasons (Barash and Arnon, 2014). The third is direct evidence of the screening gap: as early as 2006, a journalist obtained a fragment of variola (smallpox) DNA by mail order from a commercial supplier without encountering any screening, an early public demonstration that the commercial synthesis supply chain could be used to acquire regulated sequence (Randerson, 2006). Together with the poliovirus (2002), 1918 influenza (2005), and horsepox (2018) reconstructions discussed elsewhere in this review, these episodes establish that the dual-use risk is neither hypothetical nor confined to a single technology: it arises from benign research, from the mere publication of a sequence, and from the accessibility of the synthesis supply chain itself. The relatively small number of documented malicious incidents to date is itself notable; the concern is primarily anticipatory, resting on the convergence of capability, falling cost, and rising accessibility rather than on a record of past attacks.

6.2 Gain-of-function research and the H5N1 controversy

A related but analytically distinct controversy concerns gain-of-function (GOF) research: experiments that deliberately enhance a pathogen's transmissibility, virulence, or immune evasion in order to study its pandemic potential. The defining case remains a pair of studies published in 2012 in which two independent laboratories, one at Erasmus Medical Center in the Netherlands and one working across the University of Wisconsin and the University of Tokyo, separately showed that highly pathogenic avian influenza A/H5N1 virus, engineered by site-directed mutagenesis and adapted through serial passage in ferrets, could acquire the ability to transmit through the air between mammals (Herfst et al., 2012; Imai et al., 2012) (Figure 3). Because ferrets are a standard model for studying respiratory transmission of influenza in humans, the results were widely interpreted as demonstrating that H5N1, a virus that at the time had not achieved efficient human-to-human transmission in nature, was potentially only a small number of mutations away from doing so (Herfst et al., 2012). When word of the unpublished findings spread in late 2011, the United States National Science Advisory Board for Biosecurity (NSABB) took the unprecedented step of recommending that key methodological details be withheld from publication, out of concern that the information could be misused; after further review and clarification by the study authors, the NSABB reversed itself in March 2012 and both studies were ultimately published in full (Fouchier et al., 2012; Herfst et al., 2012). The episode prompted the NSABB's own multi-year deliberative process on how such research should be

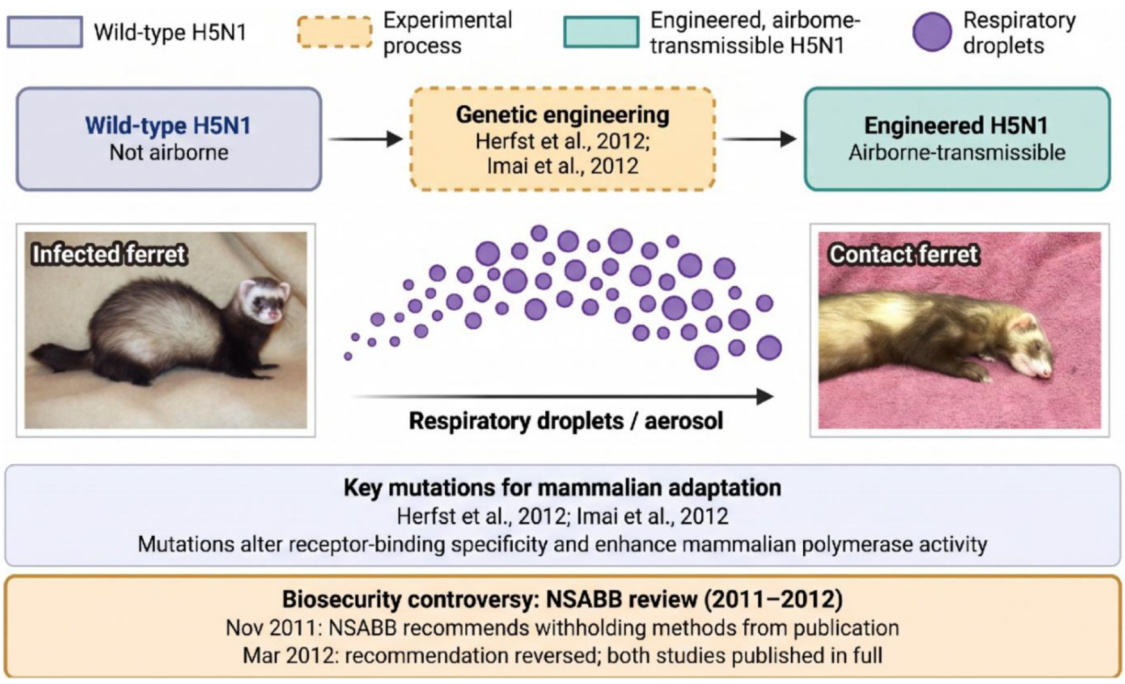


FIGURE 3
Gain-of-function experiments with H5N1 influenza virus (Herfst et al. and Imai et al., 2012). Wild-type H5N1 (not airborne in mammals) was subjected to site-directed mutagenesis and serial passage in ferrets independently by two groups. The resulting engineered variants transmitted by respiratory droplets between ferrets. Key mutations span the haemagglutinin (HA: N182K, Q222L, G224S), RNA polymerase subunits (PB2: E627K; PB1: K391E, E581G, A661T), and the acidic polymerase subunit (PA: K356R), collectively shifting receptor-binding specificity toward human-type sialic acids and enhancing mammalian polymerase activity. The findings triggered an unprecedented NSABB review: the board initially recommended withholding key methodological details from publication (November 2011) before reversing its decision in March 2012; both studies were ultimately published in full (Fouchier et al., 2012; Herfst et al., 2012; Imai et al., 2012).

evaluated and overseen going forward, culminating in a formal 2016 report recommending a structured risk-benefit review process for proposed gain-of-function studies with pandemic-potential pathogens (National Science Advisory Board for Biosecurity, 2016), which in turn shaped the government-wide oversight framework discussed in Section 7.1.

The origin of SARS-CoV-2 — whether a natural zoonotic spillover or a research-associated event—remains debated and unresolved, and we take no position on it here; the question has been examined extensively elsewhere. We mention it because, independent of its eventual resolution, the plausibility of a research-associated origin for a pandemic of this scale has sharply intensified scrutiny of gain-of-function research and laboratory biosafety and has helped drive the governance developments discussed in Section 7.

Gain-of-function experiments are not the only, and often not the safest, route to understanding viral determinants of transmissibility or virulence. Loss-of-function approaches, in which candidate virulence or host-range genes are knocked out or knocked down by targeted mutagenesis, RNA interference, or CRISPR-based editing, can identify the genetic basis of a phenotype without generating a novel pathogen of enhanced capability. Reverse-genetics systems can likewise be used to attenuate rather than enhance, and replication-defective or single-cycle viral systems, pseudotyped particles, and chimeric surrogate viruses allow entry, receptor usage, and antigenicity to be studied at markedly lower biosafety risk because no transmissible, fully competent pathogen is produced. Increasingly, computational prediction of phenotype from sequence, and structure-based inference of receptor binding, offer a route to anticipating the consequences of candidate mutations *in silico* before, or instead of, constructing them. These complementary strategies do not render gain-of-function studies categorically unnecessary, but they narrow the set of experiments for which deliberately enhancing a pathogen is the only informative design, which is precisely the set that the risk-benefit review processes discussed in Section 7 are intended to scrutinize.

6.3 Falling costs, rising access

A structural trend underlying both the dual-use and gain-of-function debates is the steadily falling cost of DNA synthesis. Cost data compiled alongside the 2018 horsepox report show a roughly exponential decline in the price of custom gene synthesis over the 2 decades preceding that study, with large clone assembly (greater than ten kilobase pairs) remaining somewhat more expensive due to the additional labor and quality control required, but following the same downward trajectory (Noyce et al., 2018). Figure 4 compiles publicly reported gene-synthesis prices from 1999 to the mid-2020s: the cost per base fell from roughly US$25 to a few US cents, a decline of more than two orders of magnitude in about 25 years (Carlson, 2009; Diggans and Leproust, 2019; Grinstein, 2023).

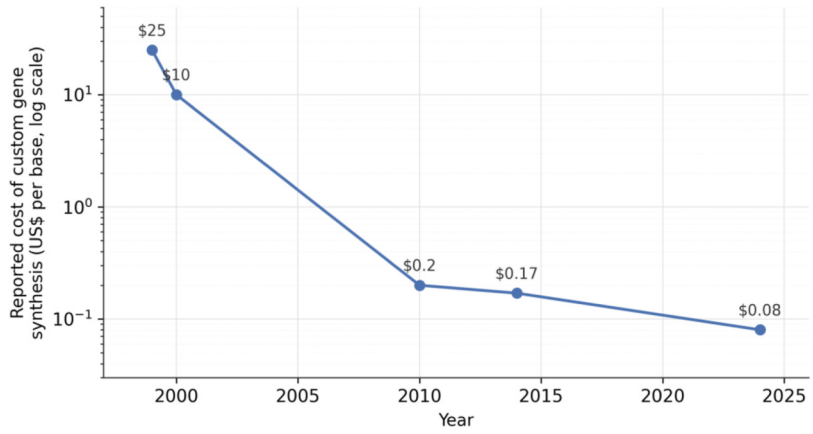


FIGURE 4
Reported cost of custom gene synthesis, 1999 to the mid-2020s. Data points are illustrative values compiled from published industry and academic sources rather than a single continuous series; the qualitative trend, a decline of roughly two to three orders of magnitude over 25 years, is consistent across sources(Carlson, 2009; Noyce et al., 2018; Diggans and Leproust, 2019; Grinstein, 2023).

That trend has continued: by the mid-2020s, dozens of commercial providers worldwide offer gene and genome synthesis services, and benchtop DNA synthesizers capable of producing genetic material outside of centralized, screenable supply chains have begun to enter the market. As early as 2019, synthesis providers themselves were publicly flagging that increasing scale and expanding synthetic biology capability would require revisiting existing screening practices to keep pace (Diggans and Leproust, 2019). Falling cost and rising accessibility do not, by themselves, create new technical capabilities beyond those already demonstrated in the milestones reviewed in Section 4, but they substantially widen the population of actors, state and non-state, well-resourced and poorly resourced, who could in principle attempt to reproduce them, which is precisely why synthesis screening (Section 7.2) has become a central focus of biosecurity policy.

Decentralized and benchtop DNA synthesis. The screening controls described above operate at the level of centralized commercial providers, which can be audited, held to common standards, and required to screen both sequences and customers. A structural shift now threatens to bypass this control point. Benchtop DNA synthesizers, desktop devices that allow a laboratory to produce synthetic DNA in-house, move synthesis out of the auditable commercial supply chain and into private settings where orders need not pass through any centralized screen. The same devices that give legitimate laboratories faster, cheaper, and more confidential access to custom DNA also, in principle, allow a user to obtain regulated sequence without external review, a governance loophole that order-level screening at commercial providers cannot close. Addressing it requires moving screening onto the device itself, so that the synthesizer evaluates a sequence before producing it, bound to the same curated databases and update mechanisms used by the Common Mechanism and comparable tools. Recent work has begun to test how distributed screening, including AI-assisted verification of customer legitimacy, might be implemented and evaluated at scale (Acelas et al., 2026). Extending mandatory, standardized screening from centralized providers to benchtop and distributed devices is, in our assessment, among the most urgent near-term priorities for synthesis-screening policy.

6.4 Artificial intelligence and the screening arms race

The most recent complication is the use of artificial intelligence to design biological sequences. Deep generative models can now produce entirely novel protein sequences and structures with specified functions, a capability with substantial legitimate value for vaccine antigen design, novel binders against viral surface proteins, and industrial biotechnology (Watson et al., 2023; Baker and Church, 2024). Because most existing biosecurity screening software works by comparing an ordered DNA sequence against databases of known pathogens and toxins using sequence similarity (homology), it is well suited to catching orders for a known dangerous sequence or a close variant of one, but may be poorly suited to catching a sequence that has been substantially redesigned by an AI tool while retaining the same three-dimensional structure and biological function (Baker and Church, 2024; Wittmann et al., 2025). In 2025, a multi-institution collaboration formally demonstrated this vulnerability: using open-source AI protein-design software, the researchers generated variants of proteins of concern that could not be reliably detected by the biosecurity screening tools then in routine use by nucleic acid synthesis providers (Wittmann et al., 2025). Rather than simply publicizing the vulnerability, the same team worked with screening providers to develop and deploy patches that substantially improved detection of these AI-redesigned, or "synthetic homolog," sequences before the underlying paper was published, an approach intended to model responsible disclosure for dual-use findings in the AI era (Wheeler et al., 2024; Wittmann et al., 2025). The episode illustrates a pattern likely to recur: as computational design tools become more capable, biosecurity screening infrastructure must be continually updated to keep pace, and the two capabilities, design and detection, are now

developing in an explicit, ongoing arms-race dynamic (Baker and Church, 2024). The 2026 demonstration of complete generative bacteriophage genome design (Sections 3.4 and 5.4) sharpens this dynamic rather than resolving it. The demonstration itself targeted a well-characterized, non-pathogenic model phage under conditions widely regarded as low risk, and the accompanying Perspective published alongside the study argued that the appropriate response is proportionate, model-organism-specific caution rather than alarm about this particular result (Inglesby and Hanke, 2026; King et al., 2026). At the same time, both the study's authors and outside commentators have noted that the underlying method, using a genome language model to design an entire functional viral genome rather than a single gene or protein, is not intrinsically restricted to bacteriophages, and that most existing screening infrastructure, including the sequence-homology approaches described above, was built to catch suspicious individual genes or short regulated sequences rather than to evaluate a fully AI-designed genome as a whole (Inglesby and Hanke, 2026; King et al., 2026). That gap is a natural extension of the AI-driven screening challenge already identified for individual proteins in Section 6.4, and reinforces the case, developed further in Section 9, for screening tools that evaluate structural and functional plausibility at the genome scale rather than relying on sequence similarity to known pathogens alone.

It is important not to overstate the immediate dual-use implications of the 2026 generative bacteriophage result. Bacteriophages and the viruses that infect humans and other mammals differ fundamentally in ways that bear directly on risk. Phages infect bacteria and are, by host range, incapable of infecting human or animal cells; the demonstration used a well-characterized, non-pathogenic model phage of *E. coli*. The recovery of a functional phage is also technically far simpler than that of most mammalian viruses: an assembled phage genome can be introduced directly into a permissive bacterial host, whereas recovering an infectious mammalian virus from synthetic material typically requires specialized reverse-genetics systems, helper functions, and permissive cell lines that are themselves significant technical barriers. The significance of the result therefore lies in its demonstration that a genome language model can design a complete, functional viral genome, a proof of concept for the generality of the method, rather than in any immediate capacity to design a human or animal pathogen (Inglesby and Hanke, 2026; King et al., 2026). This host restriction, however, should not place bacteriophages wholly outside the sphere of human-health concern: the human body is itself a holobiont, hosting a vast commensal microbiota whose bacterial members are integral to metabolism, immune development, and resistance to colonization by pathogens. Engineered or AI-designed phages directed against these beneficial taxa could in principle be used to disrupt the microbiome deliberately, a distinct and underappreciated dual-use consideration that applies even though phages cannot infect human cells directly.

A complementary 2026 analysis reinforces this point, confirming that current sequence-based screening tools systematically fail to detect AI-redesigned variants of proteins of concern and arguing for screening that incorporates structural and functional prediction rather than homology alone (Wittmann et al., 2026).

7 Governance, biosafety, and biosecurity frameworks

The risks described in Section 6 have not gone unaddressed. A patchwork of institutional, national, and international mechanisms has developed specifically to govern dual-use life-sciences research and to reduce the chance that synthetic virology capabilities are misused, even as most observers agree that the current architecture remains incomplete and unevenly enforced.

7.1 Institutional and national oversight

In the United States, the NSABB advises federal agencies on the identification and management of dual-use research of concern and played a central role in the 2011–2012 H5N1 controversy described in Section 6.2, culminating in its 2016 report on the evaluation and oversight of proposed gain-of-function research (Fouchier et al., 2012; 2012; National Science Advisory Board for Biosecurity, 2016). Following that report, the United States government developed the P3CO (Potential Pandemic Pathogen Care and Oversight) framework, an enhanced, multi-agency review process specifically for federally funded research reasonably anticipated to create, transfer, or use enhanced potential pandemic pathogens, requiring institutional and federal-level risk-benefit review before such research can (Office of Science and Technology Policy, 2017). Other countries and institutions maintain analogous, though not always harmonized, biosafety and biosecurity committees, export-control regimes, and funder-level review requirements. The European Union illustrates both this pattern of non-harmonized national arrangements and the broader lag between technical capability and dedicated oversight described throughout this section. EU-level control of dual-use life-sciences work has historically relied mainly on export-control instruments, principally the Dual-Use Regulation (EU) 2021/821, which requires authorization for the cross-border transfer of a list of items and technologies with both civilian and security applications, rather than on a dedicated mechanism for reviewing or screening dual-use research (Regulation, 2021/821 of the European Parliament and of the Council of 20 Ma y 2021 setting up a Union regime for the control of exports, brokering, technical assistance, transit and transfer of dual-use items (recast), 2021). Contained use of genetically modified organisms is regulated at the laboratory level under Directive 2009/41/EC, transposed with some variation across Member States, but until recently the Union had no binding, harmonized requirement that commercial DNA and RNA synthesis providers screen customer orders against databases of pathogens of concern, leaving that function to voluntary industry self-regulation such as the International Gene Synthesis Consortium (Directive, 2009/41/EC of the European Parliament and of the Council of 6 Ma y 2009 on the contained use of genetically modified micro-organisms, 2009; Diggans and Leproust, 2019). This gap began to close only in December 2025, when the European Commission proposed the European Biotech Act, which would introduce the Union's first dedicated biosecurity

provisions, including a biodefence chapter and an independent Advisory Group on Biosecurity tasked with monitoring risks from AI models applied to (European Commission, 2025). As of this writing, the proposal remains under negotiation between the European Parliament and the Council, with adoption not expected before the end of 2026, a further instance of governance following rather than anticipating technical advance.

Outside North America and the European Union, national approaches to DURC oversight vary widely in both instrument and enforcement. China moved earliest and most comprehensively among non-Western jurisdictions, enacting its first Biosecurity Law in October 2020 (effective April 2021, amended 2024), which established a dedicated chapter on the security of biotechnology research, development, and application, required government approval before conducting research on pathogens that China has declared eliminated or not yet found within its territory, and introduced a hierarchical, risk-based licensing system for pathogen laboratories (Standing Committee of the National People's Congress of the People's Republic of China, 2020; Cao, 2021). China has also taken an active role in the international, voluntary dimension of DURC governance: together with Pakistan and Tianjin University, and in collaboration with the InterAcademy Partnership and the Johns Hopkins Center for Health Security, Chinese institutions co-developed the Tianjin Biosecurity Guidelines for Codes of Conduct for Scientists, a set of ten non-binding principles submitted to the Biological Weapons Convention's Ninth Review Conference and endorsed by the InterAcademy Partnership as a template for institutional and national codes of conduct (Tianjin Biosecurity Guidelines for Codes of Conduct for Scientists, 2021). India, by contrast, continues to rely on legacy instruments, principally the Department of Biotechnology's recombinant DNA guidelines (first issued in 1989 and periodically revised) and the Biological Diversity Act of 2002, neither of which was substantially amended in response to COVID-19 or extended to require dedicated screening of commercial gene-synthesis orders, leaving DURC oversight fragmented across ministries rather than centralized under a single biosecurity authority (International Bar Association, 2025).

Canada regulates human pathogens and toxins nationally through the Human Pathogens and Toxins Act (HPTA), administered by the Public Health Agency of Canada, which requires licensed facilities to maintain a Plan for Administrative Oversight that explicitly addresses dual-use risk. The Act's practical limits were illustrated in 2019, when material including Ebola and Nipah virus samples was exported from Canada's National Microbiology Laboratory to the Wuhan Institute of Virology under internal institutional approvals, an episode that received wider scrutiny once related correspondence was later declassified by the Canadian Parliament; amendments to the HPTA intended to strengthen biosecurity oversight of scientific research received Royal Assent in March 2026 (Human Pathogens and Toxins Act, S.C., 2009). Across the Middle East, dual-use governance remains comparatively heterogeneous: Israel has developed one of the region's most structured dual-use review frameworks, prompted in part by the same 2011–2012 H5N1 controversy discussed in Section 6.2, whereas most other states in the region rely primarily on Biological Weapons Convention membership and multilateral export-control arrangements such as the Australia Group rather than on dedicated domestic DURC legislation, a gap that remains a priority for future policy development in the region (Haddad, 2025).

7.2 Nucleic acid synthesis screening

Because commercial gene synthesis is the practical bottleneck for building most viral genomes larger than a few genes, screening synthesis orders against databases of regulated pathogens and toxins has become a central biosecurity control point. Many, though historically not all, commercial DNA synthesis providers voluntarily screen customer orders, both the ordered sequences and the identity of the customer, following industry best-practice guidance developed collaboratively over the past decade, but until recently no country legally required such screening, and no single technical standard existed across providers; the practical consequence of that absence had already been demonstrated publicly as early as 2006 (Randerson, 2006; Diggans and Leproust, 2019). The International Biosecurity and Biosafety Initiative for Science (IBBIS), an independent organization headquartered in Geneva, Switzerland, was established to address this gap; its flagship project, the Common Mechanism, is a free, open-source software package that screens nucleic acid sequences of fifty base pairs or more against curated databases of biorisk sequences and regulated pathogens, developed in consultation with an international technical consortium of synthesis providers, biosecurity specialists, and policymakers (Wheeler et al., 2024; International Biosecurity and Biosafety Initiative for Science, 2026). As described in Section 6.4, the Common Mechanism and comparable tools are now being actively hardened against AI-generated sequence variants designed to evade homology-based detection, an effort that is expected to require continual updates as generative design tools improve (Wheeler et al., 2024; Wittmann et al., 2025).

7.3 International coordination

Above the level of individual institutions and countries, coordination occurs through instruments such as the Biological Weapons Convention, which prohibits the development, production, acquisition, transfer, and stockpiling of biological and toxin weapons and which, as of the mid-2020s, has been joined by 189 states (United Nations Office for Disarmament Affairs UNODA, 1972), and through multilateral export-control arrangements that restrict trade in certain pathogens, toxins, and dual-use equipment.

These instruments predate the synthetic biology era and were not designed with gene synthesis or AI-assisted design specifically in mind, which is part of the motivation behind newer, technology-specific efforts such as the Common Mechanism. A parallel strand of the biosecurity literature argues that the problem is partly definitional: the categories on which the Convention and allied export-control regimes rest were framed around agents and stockpiles rather than around design capability, and may need to be reconceptualized to encompass AI- and CRISPR-enabled routes to a biological weapon (Haddad, 2025). A recurring theme in the biosecurity literature is that governance instruments generally lag the pace

of technical advance: the P3CO framework followed, rather than anticipated, the 2012 H5N1 controversy, and synthesis-screening standardization efforts followed, rather than preceded, the proliferation of low-cost commercial gene synthesis (Office of Science and Technology Policy, 2017; Wheeler et al., 2024; International Biosecurity and Biosafety Initiative for Science, 2026).

8 Ethical debates: Who decides how far is too far?

Beneath the specific policy mechanisms described in Section 7 lies a set of unresolved ethical questions that recur across the cases discussed throughout this review.

8.1 Publication and the information-hazard dilemma

Both the H5N1 gain-of-function studies and the horsepox synthesis forced journals, funders, and governments to confront the same underlying dilemma: withholding methodological detail from a published paper may reduce the risk of misuse but also limits the ability of the broader scientific and biosecurity community to understand, prepare for, and defend against the underlying risk, while full publication accelerates legitimate science but also lowers the barrier to reproduction by a less well-intentioned actor (Fouchier et al., 2012; Koblentz, 2017; Coyne, 2018; Noyce and Evans, 2018). Some biosecurity researchers have argued for structural changes to how such findings are reviewed and released in the first place, for instance by separating the most sensitive methodological details from the main publication or by building information-hazard assessment directly into the peer-review process, rather than treating each disclosure decision as a one-off editorial judgment call (Esvelt, 2018). There is no settled consensus on how to resolve this trade-off in general; existing frameworks such as NSABB review and journal-level dual-use policies handle it case by case, and the 2025 approach of quietly hardening screening defenses before publicizing a newly discovered AI-related vulnerability represents one attempt at a middle path specific to the current technological moment (Wheeler et al., 2024; Wittmann et al., 2025).

8.2 Who bears the risk, and who reaps the benefit

A second recurring tension is distributive: the benefits of synthetic virology, improved vaccines, cancer therapies, and gene therapies, accrue disproportionately to whoever can access and afford them, while the risks of misuse are borne globally and are largely outside any individual researcher's or country's control. This asymmetry underlies calls for international, rather than purely national, coordination on synthesis screening and dual-use oversight, on the grounds that a dangerous synthetic pathogen released anywhere could threaten populations everywhere, regardless of where the underlying research or synthesis occurred (United Nations Office for Disarmament Affairs UNODA, 1972; Wheeler et al., 2024; International Biosecurity and Biosafety Initiative for Science, 2026).

8.3 Inevitability versus restraint

Finally, the field has not resolved a more basic disagreement about whether restraint is even an effective strategy. One view, expressed explicitly by the horsepox researchers themselves, holds that because the underlying technology is generalizable, accessible, and improving on its own trajectory, attempting to prevent any given synthesis milestone from being achieved is largely futile, and that resources are better spent on defensive capabilities such as surveillance, rapid vaccine platforms, and synthesis screening (Noyce and Evans, 2018). The opposing view, developed most fully in formal ethical analyses of gain-of-function research, holds that specific high-consequence experiments, particularly those involving eradicated pathogens or enhanced pandemic potential, can and should be delayed or redirected through review processes such as P3CO and NSABB oversight, drawing an explicit analogy to research-ethics principles that require even scientifically valuable experiments to be foregone when their risk to third parties is disproportionate to their benefit and safer alternatives exist (Lipsitch and Galvani, 2014; Selgelid, 2016). Both positions are represented among credentialed biosecurity and bioethics experts, and the tension between them is unlikely to be resolved by technical means alone.

9 Future outlook: How far can we go?

Three trends identified throughout this review are likely to jointly define the field's near-term trajectory: continued growth in the size and complexity of genomes that can be synthesized and designed; continued decline in the cost and continued increase in the accessibility of DNA synthesis; and continued improvement in AI-assisted sequence and protein design, which both expands legitimate capabilities and complicates biosecurity screening (Hutchison et al., 2016; Noyce et al., 2018; Watson et al., 2023; Baker and Church, 2024; Wittmann et al., 2025). None of these trends shows signs of reversing, which suggests that the technical answer to "how far can we go" is, for the foreseeable future, "farther than we currently can," in both the beneficial and the concerning directions simultaneously.

On the beneficial side, the convergence of computational design with synthetic virology points toward vaccines and therapeutics that can be designed and produced on timescales measured in weeks rather than years, self-amplifying and circular RNA platforms that further reduce dose and cost (Bloom et al., 2021; Casmil et al., 2025), and oncolytic and gene-therapy vectors engineered with increasing precision for specific tumor types or genetic disorders using a combination of rational design and directed evolution (Andtbacka et al., 2015; Nobel Prize in Nobel Prize in Chemistry, 2018; High and Roncarolo, 2019; Kaufman, 2022). Realizing this potential responsibly will likely require continued investment in platforms that generalize across pathogens, so that pandemic preparedness does not depend on having anticipated the exact virus that eventually emerges.

On the concerning side, the central lesson of Section 6 is that the bottleneck on misuse is shifting from raw technical capability, which the milestones in Section 4 show are already extensive, toward the effectiveness of screening, oversight, and international coordination. This implies that future governance will need to be adaptive rather than fixed: static rules written for a particular technology, such as homology-based sequence screening, are likely to be outpaced by the next capability, such as generative protein design, unless review processes are built with the expectation of continual revision (Wheeler et al., 2024; Wittmann et al., 2025). Several concrete directions follow from the evidence reviewed here: extending mandatory, standardized synthesis screening to benchtop and distributed DNA synthesis devices, not just centralized commercial providers (Diggans and Leproust, 2019; Wheeler et al., 2024; International Biosecurity and Biosafety Initiative for Science, 2026); building screening tools that combine sequence homology with structural and functional prediction, so that AI-redesigned sequences of concern remain detectable even when they no longer closely resemble known pathogens at the sequence level (Baker and Church, 2024; Wittmann et al., 2025); institutionalizing information-hazard assessment earlier in the research and publication pipeline rather than treating it as an after-the-fact editorial judgment (Esvelt, 2018); and maintaining international coordination mechanisms capable of updating biosecurity standards at the pace of the underlying technology, rather than years afterward, as has historically been the case (United Nations Office for Disarmament Affairs UNODA, 1972; Office of Science and Technology Policy, 2017; Wheeler et al., 2024; International Biosecurity and Biosafety Initiative for Science, 2026).

Beyond governance, several scientific strategies can reduce risk at the level of experimental design itself. Engineered biocontainment, including synthetic auxotrophy or dependence on non-natural amino acids, genetic kill-switches, and split or trans-complementation systems that require a dedicated complementing cell line, can help ensure that an engineered construct cannot propagate outside the laboratory. Such safer-by-design measures complement, rather than replace, the loss-of-function alternatives discussed in Section 6.2.

On the normative side, the heterogeneity of standards across countries and across academic, industrial, and government laboratories is itself a source of risk, because the overall level of protection is set by the least stringent accessible node. Convergence toward a common, broadly applicable best-practice baseline is therefore a priority, and several building blocks already exist that could be harmonized: the International Gene Synthesis Consortium's Harmonized Screening Protocol and the free, open-source IBBIS Common Mechanism provide a shared technical standard for synthesis screening; ISO 35001 provides a certifiable biorisk-management standard applicable to any laboratory regardless of sector; the WHO Global Guidance Framework for the Responsible Use of the Life Sciences offers a neutral, internationally negotiated set of values and tools; and the Tianjin Biosecurity Guidelines offer a template code of conduct for individual scientists (International Organization for Standardization, 2019; World Health Organization, 2022). A consistent, auditable baseline assembled from these instruments and adopted across sectors and jurisdictions, rather than voluntarily and unevenly, would do more to raise the global floor of safe practice than any single national rule.

The honest answer to the question posed in this review's title is therefore not a fixed technical ceiling but a moving one, set jointly by what biology and computation make possible and by what oversight, screening, and international norms are able to catch and constrain. The history reviewed here suggests that the technical Frontier will keep advancing; whether governance keeps pace with it is, at this point, more a matter of policy and institutional will than of biology.

10 Conclusion

Synthetic virology has advanced from the first chemical synthesis of an infectious poliovirus genome in 2002 to the assembly of a 212-kilobase-pair synthetic orthopoxvirus in 2018, a roughly thirty-fold increase in achievable genome complexity in 16 years, driven by improvements in DNA synthesis, genome assembly, directed evolution, and, most recently, AI-assisted sequence design, culminating in 2026 in the first generative design of a complete, functional bacteriophage genome (Cello et al., 2002; Hutchison et al., 2016; Noyce et al., 2018; Wittmann et al., 2025; King et al., 2026). The same toolkit has delivered substantial public health benefit, including mRNA and self-amplifying RNA vaccine platforms, approved oncolytic virus therapies, engineered gene-therapy vectors, and a revival of phage therapy against antibiotic-resistant bacteria (Andtbacka et al., 2015; High and Roncarolo, 2019; Kortright et al., 2019; Polack et al., 2020; Hồ et al., 2024). It has also produced the field's defining biosecurity concerns: the dual-use research of concern debate crystallized by the horsepox synthesis, the gain-of-function controversy surrounding transmissible H5N1 influenza, the steadily falling cost and rising accessibility of gene synthesis, and the newly emerging risk that AI-designed sequences may evade existing biosecurity screening (Herfst et al., 2012; Imai et al., 2012; DiEuliis et al., 2017; Koblentz, 2017; Noyce and Evans, 2018; Noyce et al., 2018; Wittmann et al., 2025). Governance mechanisms, including NSABB review, the P3CO framework, and international nucleic acid synthesis screening initiatives such as the Common Mechanism, have developed in response to these concerns but have consistently lagged behind the technology they are meant to govern (National Science Advisory Board for Biosecurity, 2016; Office of Science and Technology Policy, 2017; Wheeler et al., 2024; International Biosecurity and Biosafety Initiative for Science, 2026). How far synthetic virology can go is, on present evidence, further than most non-specialists assume; how far it should go, and how quickly oversight can be done to keep pace with capability, remain open and consequential questions for the scientific, biosecurity, and policy communities alike.

Author contributions

TC: Data curation, Methodology, Writing – review and editing, Investigation. LR: Funding acquisition, Writing – review and editing, Supervision, Resources. DC: Writing – original draft, Conceptualization, Writing – review and editing.

Funding

The author(s) declared that financial support was received for this work and/or its publication. This work was supported by the European Union's HORIZON-EIC-2022-PATHFINDEROPEN-01 (GA 101098972): "Bacteria Biofilm as bio-factory for tissue regeneration" -BIOACTION.

Acknowledgments

The authors thank the Department of Health Sciences, Università del Piemonte Orientale, for its institutional support. T.V.C. gratefully acknowledges funding from the Ministry of University and Research (MUR) and from the European Union's Horizon Europe research and innovation programme under grant agreement no. 101098972 (BIOACTION) for his PhD scholarship. This manuscript originated from the Synthetic Biology course (MS in Medical Biotechnology), taught by DC

Conflict of interest

The author(s) declared that this work was conducted in the absence of any commercial or financial relationships that could be construed as a potential conflict of interest.

Generative AI statement

The author(s) declared that generative AI was used in the creation of this manuscript. Claude Sonnet five and Claude Opus 4.8 (Anthropic) were used to assist with literature search, drafting and revising text, generating tables and figures from data compiled from cited sources, and formatting. All AI-assisted content, including factual claims, citations, and data underlying tables and figures, was verified by the authors against primary sources. The authors take full responsibility for the accuracy and integrity of the content of this publication.

Any alternative text (alt text) provided alongside figures in this article has been generated by Frontiers with the support of artificial intelligence and reasonable efforts have been made to ensure accuracy, including review by the authors wherever possible. If you identify any issues, please contact us.

Publisher's note

All claims expressed in this article are solely those of the authors and do not necessarily represent those of their affiliated organizations, or those of the publisher, the editors and the reviewers. Any product that may be evaluated in this article, or claim that may be made by its manufacturer, is not guaranteed or endorsed by the publisher.

References

Acelas, A., Palya, H., Flyangolts, K., Fady, P.-E., and Nelson, C. (2026). Evaluating AI-assisted customer verification for synthetic nucleic acid screening. *Front. Bioeng. Biotechnol.* 14, 1819556. doi: 10.3389/fbioe.2026.1819556

Amgen (2015). *FDA Approves IMLYGIC (Talimogene Laherparepvec) as First Oncolytic Viral Therapy in the US.* Thousand Oaks, CA, USA: Amgen.

Andtbacka, R. H. I., Kaufman, H. L., Collichio, F., Amatruda, T., Senzer, N., Chesney, J., et al. (2015). Talimogene laherparepvec improves durable response rate in patients with advanced melanoma. *J. Clin. Oncol.* 33, 2780–2788. doi: 10.1200/JCO.2014.58.3377

Baden, L. R., El Sahly, H. M., Essink, B., Kotloff, K., Frey, S., Novak, R., et al. (2021). Efficacy and safety of the mRNA-1273 SARS-CoV-2 vaccine. *N. Engl. J. Med.* 384, 403–416. doi: 10.1056/nejmoa2035389

Baker, D., and Church, G. (2024). Protein design meets biosecurity. *Science* 383, 349–350. doi: 10.1126/science.ado1671

Barash, J. R., and Arnon, S. S. (2014). A novel strain of Clostridium botulinum that produces type B and type H botulinum toxins. *J. Infect. Dis.* 209, 183–191. doi: 10.1093/infdis/jit449

Bloom, K., van den Berg, F., and Arbuthnot, P. (2021). Self-amplifying RNA vaccines for infectious diseases. *Gene Ther.* 28, 117–129.

Cao, C. (2021). China's evolving biosafety/biosecurity legislations. *J. Law Biosci.* 8, lsab018. doi: 10.1093/jlb/lsab020

Carlson, R. (2009). The changing economics of DNA synthesis. *Nat. Biotechnol.* 27, 1091–1094. doi: 10.1038/nbt1209-1091

Casmil, I. C., Jin, J., Won, E.-J., Huang, C., Liao, S., Cha-Molstad, H., et al. (2025). The advent of clinical self-amplifying RNA vaccines. *Mol. Ther.* 33, 2565–2582. doi: 10.1016/j.ymthe.2025.03.060

Cello, J., Paul, A. V., and Wimmer, E. (2002). Chemical synthesis of poliovirus cDNA: generation of infectious virus in the absence of natural template. *Science* 297, 1016–1018. doi: 10.1126/science.1072266

Corbett, K. S., Edwards, D. K., Leist, S. R., Abiona, O. M., Boyoglu-Barnum, S., Gillespie, R. A., et al. (2020). SARS-CoV-2 mRNA vaccine design enabled by prototype pathogen preparedness. *Nature* 586, 567–571. doi: 10.1038/s41586-020-2622-0

Coyne, C. B. (2018). Horsepox: framing a dual use research of concern debate. *PLoS Pathog.* 14, e1007344. doi: 10.1371/journal.ppat.1007344

DiEuliis, D., Berger, K., and Gronvall, G. (2017). Biosecurity implications for the synthesis of horsepox, an orthopoxvirus. *Health Secur.* 15, 629–637. doi: 10.1089/hs.2017.0081

Diggans, J., and Leproust, E. (2019). Next steps for access to safe, secure DNA synthesis. *Front. Bioeng. Biotechnol.* 7, 86. doi: 10.3389/fbioe.2019.00086

Directive (2009). *Directive 2009/41/EC of the European Parliament and of the Council of 6 May 2009 on the Contained Use of Genetically Modified Micro-organisms.*

Esvelt, K. M. (2018). Inoculating science against potential pandemics and information hazards. *PLoS Pathog.* 14, e1007286. doi: 10.1371/journal.ppat.1007286

European Commission (2025). *Proposal for a Regulation of the European Parliament and of the Council on Establishing a Framework of Measures for Strengthening Union's Biotechnology and Biomanufacturing Sectors Particularly in the Area of Health (European Biotech Act).*

Fouchier, R. A. M., Herfst, S., and Osterhaus, A. D. M. E. (2012). Restricted data on influenza H5N1 virus transmission. *Science* 335, 662–663. doi: 10.1126/science.1218376

Gibson, D. G., Benders, G. A., Andrews-Pfannkoch, C., Denisova, E. A., Baden-Tillson, H., Zaveri, J., et al. (2008). Complete chemical synthesis, assembly, and cloning of a Mycoplasma genitalium genome. *Science* 319, 1215–1220. doi: 10.1126/science.1151721

Gibson, D. G., Young, L., Chuang, R.-Y., Venter, J. C., Hutchison, C. A., and Smith, H. O. (2009). Enzymatic assembly of DNA molecules up to several hundred kilobases. *Nat. Methods* 6, 343–345. doi: 10.1038/nmeth.1318

Graham, B. S. (2020). Rapid COVID-19 vaccine development. *Science* 368, 945–946. doi: 10.1126/science.abb8923

Grinstein, J. D. (2023). "The long and winding road: on-demand DNA synthesis in high demand," in *Genetic Engineering & Biotechnology News*. Available online at: https://www.genengnews.com/topics/genome-editing/the-long-and-winding-road-on-demand-dna-synthesis-in-high-demand/(Accessed August 26, 2026).

Haddad, H. (2025). Redefining biological weapons in the evolving AI, CRISPR, and biothreat landscape. *Ethics, Med. Public Health* 33, 101176. doi: 10.1016/j.jemep.2025.101176

Herfst, S., Schrauwen, E. J. A., Linster, M., Chutinimitkul, S., de Wit, E., Munster, V. J., et al. (2012). Airborne transmission of influenza A/H5N1 virus between ferrets. *Science* 336, 1534–1541. doi: 10.1126/science.1213362

High, K. A., and Roncarolo, M. G. (2019). Gene therapy. *N. Engl. J. Med.* 381, 455–464. doi: 10.1056/NEJMra1706910

Hồ, N. T., Hughes, S. G., Ta, V. T., Phan, L. T., Đỗ, Q., Nguyễn, T. V., et al. (2024). Safety, immunogenicity and efficacy of the self-amplifying mRNA ARCT-154 COVID-19 vaccine: pooled phase 1, 2, 3a and 3b randomized, controlled trials. *Nat. Commun.* 15, 4081. doi: 10.1038/s41467-024-47905-1

Human Pathogens and Toxins Act, S.C. (2009). *Human Pathogens and Toxins Act, S.C.* Canada, 24.

Hutchison, C. A., Chuang, R. Y., Noskov, V. N., Assad-Garcia, N., Deerinck, T. J., Ellisman, M. H., et al. (2016). Design and synthesis of a minimal bacterial genome. *Science* 351, aad6253. doi: 10.1126/science.aad6253

Imai, M., Watanabe, T., Hatta, M., Das, S. C., Ozawa, M., Shinya, K., et al. (2012). Experimental adaptation of an influenza H5 HA confers respiratory droplet transmission to a reassortant H5 HA/H1N1 virus in ferrets. *Nature* 486, 420–428. doi: 10.1038/nature10831

Inglesby, T. V., and Hanke, M. S. (2026). AI-designed viral genomes. *Science* 393, 563–564. doi: 10.1126/science.aej8512

International Bar Association (2025). Blurred lines: Law, science, and the regulation of dual-use biological research. Available online at: https://www.ibanet.org/dual-use-biological-research-law (Accessed August 27, 2026).

International Biosecurity and Biosafety Initiative for Science (IBBIS) (2026). The common mechanism for DNA synthesis screening. Available online at: https://ibbis.bio (Accessed August 26, 2026).

International Organization for Standardization (2019). *Biorisk Management for Laboratories and Other Related Organisations* (ISO 35001:2019). Available online at: https://www.iso.org/standard/71293.html (Accessed October 18, 2026).

Jackson, R. J., Ramsay, A. J., Christensen, C. D., Beaton, S., Hall, D. F., and Ramshaw, I. A. (2001). Expression of mouse interleukin-4 by a recombinant ectromelia virus suppresses cytolytic lymphocyte responses and overcomes genetic resistance to mousepox. *J. Virol.* 75, 1205–1210. doi: 10.1128/JVI.75.3.1205-1210.2001

Karikó, K., Buckstein, M., Ni, H., and Weissman, D. (2005). Suppression of RNA recognition by toll-like receptors: the impact of nucleoside modification and the evolutionary origin of RNA. *Immunity* 23, 165–175. doi: 10.1016/j.immuni.2005.06.008

Kaufman, H. L., Shalhout, S. Z., and Iodice, G. (2022). Talimogene laherparepvec: moving from first-in-class to best-in-class. *Front. Mol. Biosci.* 9, 834841. doi: 10.3389/fmolb.2022.834841

King, S. H., Driscoll, C. L., Li, D. B., Guo, D., Merchant, A. T., Brixi, G., et al. (2026). Generative design of bacteriophages with genome language models. *Science* 393, eaec2657. doi: 10.1126/science.aec2657

Koblentz, G. D. (2017). The *de novo* synthesis of horsepox virus: implications for biosecurity and recommendations for preventing the reemergence of smallpox. *Health Secur.* 15, 620–628. doi: 10.1089/hs.2017.0061

Kortright, K. E., Chan, B. K., Koff, J. L., and Turner, P. E. (2019). Phage therapy: a renewed approach to combat antibiotic-resistant bacteria. *Cell Host & Microbe* 25, 219–232. doi: 10.1016/j.chom.2019.01.014

Lipsitch, M., and Galvani, A. P. (2014). Ethical alternatives to experiments with novel potential pandemic pathogens. *PLoS Med.* 11, e1001646. doi: 10.1371/journal.pmed.1001646

National Science Advisory Board for Biosecurity (NSABB) (2016). "Recommendations for the evaluation and oversight of proposed gain-of-function research,". Bethesda, MD, USA. *Natl. Institutes Health, Office Sci. Policy*. Available online at: https://osp.od.nih.gov/wp-content/uploads/2016/06/NSABB_Final_Report_Recommendations_Evaluation_Oversight_Proposed_Gain_of_Function_Research.pdf.

Nobel Prize in Chemistry (2018). *Nature*. Available online at: https://www.nobelprize.org/prizes/chemistry/2018/summary/ (Accessed September 1, 2026).

Nobel Prize in Physiology or Medicine (2023). *Nature*. Available online at: https://www.nobelprize.org/prizes/medicine/2023/summary/ (Accessed September 1, 2026).

Noyce, R. S., and Evans, D. H. (2018). Synthetic horsepox viruses and the continuing debate about dual use research. *PLoS Pathog.* 14, e1007025. doi: 10.1371/journal.ppat.1007025

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

Office of Science and Technology Policy (OSTP) (2017). *Recommended Policy Guidance for Departmental Development of Review Mechanisms for Potential Pandemic Pathogen Care and Oversight (P3CO)*. Washington, DC, USA: Executive Office of the President.

Pizza, M., Scarlato, V., Masignani, V., Giuliani, M. M., Arico, B., Comanducci, M., et al. (2000). Identification of vaccine candidates against serogroup B meningococcus by whole-genome sequencing. *Science* 287, 1816–1820. doi: 10.1126/science.287.5459.1816

Polack, F. P., Thomas, S. J., Kitchin, N., Absalon, J., Gurtman, A., Lockhart, S., et al. (2020). Safety and efficacy of the BNT162b2 mRNA Covid-19 vaccine. *N. Engl. J. Med.* 383, 2603–2615. doi: 10.1056/NEJMoa2034577

Racaniello, V. R., and Baltimore, D. (1981). Molecular cloning of poliovirus cDNA and determination of the complete nucleotide sequence of the viral genome. *Proc. Natl. Acad. Sci. U. S. A.* 78, 4887–4891. doi: 10.1073/pnas.78.8.4887

Randerson, J. (2006). The lax laws could allow assembly of deadly virus DNA. *Guardian*. Available online at: https://www.theguardian.com/world/2006/jun/14/terrorism.topstories3 (Accessed October 25, 2026).

Rappuoli, R. (2000). Reverse vaccinology. *Curr. Opin. Microbiol.* 3, 445–450. doi: 10.1016/s1369-5274(00)00119-3

Regulation (EU) (2021). *2021/821 of the European Parliament and of the Council of 20 May 2021 Setting up a Union Regime for the Control of Exports, Brokering, Technical Assistance, Transit and Transfer of dual-use Items (Recast)*.

Selgelid, M. J. (2016). Gain-of-function research: ethical analysis. *Sci. Eng. Ethics* 22, 923–964. doi: 10.1007/s11948-016-9810-1

Smith, H. O., Hutchison, C. A., Pfannkoch, C., and Venter, J. C. (2003). Generating a synthetic genome by whole genome assembly: PhiX174 bacteriophage from synthetic oligonucleotides. *Proc. Natl. Acad. Sci. U. S. A.* 100, 15440–15445. doi: 10.1073/pnas.2237126100

Standing Committee of the National People's Congress of the People's Republic of China (2020). *Biosecurity Law of the People's Republic of China*.

Tettelin, H., Saunders, N. J., Heidelberg, J., Jeffries, A. C., Nelson, K. E., Eisen, J. A., et al. (2000). Complete genome sequence of Neisseria meningitidis serogroup B strain MC58. *Science* 287, 1809–1815. doi: 10.1126/science.287.5459.1809

Tianjin Biosecurity Guidelines for Codes of Conduct for Scientists (2021). *J. Biosaf. Biosecurity* 3, 82–83.

Tumpey, T. M., Basler, C. F., Aguilar, P. V., Zeng, H., Solorzano, A., Swayne, D. E., et al. (2005). Characterization of the reconstructed 1918 Spanish influenza pandemic virus. *Science* 310, 77–80. doi: 10.1126/science.1119392

United Nations Office for Disarmament Affairs (UNODA) (1972). *Convention on the Prohibition of the Development, Production and Stockpiling of Bacteriological (Biological) and Toxin Weapons and on their Destruction (Biological Weapons Convention)*.

Watson, J. L., Juergens, D., Bennett, N. R., Trippe, B. L., Yim, J., Eisenach, H. E., et al. (2023). *De novo* design of protein structure and function with RFdiffusion. *Nature* 620, 1089–1100. doi: 10.1038/s41586-023-06415-8

Wheeler, N. E., Carter, S. R., Alexanian, T., Isaac, C., Yassif, J., and Millett, P. (2024). Developing a common global baseline for nucleic acid synthesis screening. *Appl. Biosaf.* 29, 71–78. doi: 10.1089/apb.2023.0034

Wimmer, E. (2006). The test-tube synthesis of a chemical called poliovirus: the simple synthesis of a virus has far-reaching societal implications. *EMBO Rep.* 7, S3–S9. doi: 10.1038/sj.embor.7400728

Wittmann, B. J., Alexanian, T., Bartling, C., Beal, J., Clore, A., Diggans, J., et al. (2025). Strengthening nucleic acid biosecurity screening against generative protein design tools. *Science* 390, 82–87. doi: 10.1126/science.adu8578

Wittmann, B. J., Wheeler, N. E., Murphy, S. T., Mitchell, T., Magalis, B. R., Gemler, B. T., et al. (2026). The limits of sequence-based biosecurity screening tools in the age of AI-assisted protein design. *Front. Bioeng. Biotechnol.* 14, 1858951. doi: 10.3389/fbioe.2026.1858951

World Health Organization (2022). Global guidance framework for the responsible use of the life sciences: mitigating biorisks and governing dual-use research. Available online at: https://www.who.int/publications/i/item/9789240056107 (Accessed October 18, 2026).