



OpenAI's patent pledge: a post-Moderna analysis

Gabriela Lenarczyk^{ID} and Mateo Aboy^{ID}*

Gabriela Lenarczyk is a Postdoctoral Fellow of the Centre for Advanced Studies in Bioscience Innovation Law (CeBIL) in the Faculty of Law at the University of Copenhagen, Copenhagen, Denmark.

Mateo Aboy is Director of Research in Technology and Law (Biomedical Innovation, AI & Law) in the Faculty of Law at the University of Cambridge, Cambridge, UK.

*Email: ma608@cam.ac.uk

Abstract

- Patent pledges have increasingly been used as private ordering mechanisms to facilitate innovation and collaboration while balancing the proprietary interests of patent holders. The *Moderna v Pfizer/BioNTech* decision by the UK High Court marked the first judicial analysis of the enforceability and interpretation of such pledges.
- OpenAI's patent pledge, committing to defensive patent use subject to specific behavioural conditions, raises novel interpretative issues. Their pledge diverges from traditional pledges by focusing on defensive use, conditioned upon specific behavioural triggers such as avoiding harm or legal actions against OpenAI.
- This article examines the pledge through the lens of the *Moderna* framework, analysing its scope, revocability and operational limits. The pledge's reliance on broad behavioural conditions and lack of explicit temporal markers introduce noteworthy legal and practical complexities. These elements affect its enforceability and create challenges for third-party reliance.
- Our analysis highlights the balance between strategic flexibility for patent holders and the need for clarity to foster trust among implementers. The findings underscore the balance between fostering innovation and protecting proprietary interests.

I. Introduction and legal context

Patent pledges—voluntary public commitments by a patent holder to refrain from exercising some or all of its rights in one or more patents—have emerged as private ordering tools for knowledge diffusion and technological advancement across industries. Such public commitments often limit patent enforcement (or other exploitation) and are typically directed at the public or specific market segments without direct compensation.¹ They serve various purposes: promoting innovation, fostering collaboration, addressing ethical concerns and responding to crises.²

Despite their increasing prevalence across technology, biotechnology and automotive sectors, patent pledges have occupied an uncertain legal territory, with limited jurisprudence addressing their enforceability, scope and revocability.

The decision of the High Court of England and Wales in *Pfizer Inc & Anor v Modernatx Inc* [2024] EWHC 1648 (Pat) (hereafter, *Moderna v Pfizer* or *Moderna*) marked the first time globally that a court examined whether patent pledges can legally grant consent to practice the patents they cover. The judgment established that, although such pledges can provide legal consent, they must meet specific requirements regarding precision, scope and clarity to be enforceable.

On 11 October 2024, OpenAI issued a patent pledge committing to use its patents only defensively, subject to specific conditions and exceptions.³ This pledge presents distinct interpretative challenges from those addressed in *Moderna*.⁴ Most notably, while *Moderna's* pledge centred on a temporal condition ('while the pandemic continues'), OpenAI's pledge is structured around behavioural conditions—specifically the absence of patent assertions or harm against the company or its users.

This article examines OpenAI's patent pledge through the interpretative framework established in *Moderna*. We analyse the

¹ G Lenarczyk, T Minssen and M Aboy, 'The Nature, Scope and Validity of Patent Pledges' (2024) 19 *JIPLP* 805–08.

² See, eg, JL Contreras and M Jacob (eds), *Patent Pledges: Global Perspectives on Patent Law's Private Ordering Frontier* (Edward Elgar 2017); G Lenarczyk, *Patent pledges na tle polskich instytucji prawnych, ze szczególnym uwzględnieniem licencji otwartej* [Patent Pledges in the Context of Polish Legal Institutions, with Special Emphasis on Licences of Right] (Publishing House of ILS PAS 2024).

³ Open AI, 'Our approach to Patents' <<https://openai.com/approach-to-patents/>> accessed 29 November 2024.

⁴ Moderna, 'Statement on Intellectual Property' <<https://investors.modernatx.com/Statements--Perspectives/Statements--Perspectives-Details/2021/Statement-on-Intellectual-Property/default.aspx>> accessed 29 November 2024.

nature of the pledge and assess its likely legal effect based on the High Court's reasoning. The analysis considers the pledge's defensive conditions, its exception clauses and their implications for both the period of applicability and the circumstances of revocation. While differing materially from Moderna's pledge in both scope and structure, OpenAI's pledge raises similar questions about the interaction between simple consent and conditional limitations. Our analysis aims to aid implementers and practitioners navigating defensive patent commitments, with particular attention to the complex interplay between breach of conditions, unilateral revocation and their differing temporal effects.

II. OpenAI's patent pledge

Patent pledges, as they continue to emerge across technological sectors from biotechnology to artificial intelligence, represent a spectrum of legal formality and effect.⁵ At one end are simple statements of intent not to assert patents; at the other are detailed, contract-like instruments specifying precise terms, conditions and durations.⁶ This diversity is not a weakness but rather reflects the mechanism's inherent flexibility—allowing pledgors to craft commitments that range from broad technology adoption initiatives to carefully calibrated instruments balancing commercial objectives with controlled IP rights release. While this flexibility historically created uncertainty about their legal effect,⁷ the High Court's analysis in *Moderna v Pfizer* now provides a useful framework for their interpretation by establishing that patent pledges can constitute simple consent to what would otherwise be infringing acts, even if they fall short of creating contractual obligations or federal law waivers.⁸ The key question is no longer whether patent pledges can create legally binding commitments, but rather what specific scope and effect a given pledge achieves through its particular formulation. OpenAI's pledge offers an instructive case study for applying this framework, beginning with its foundational declarations and operational provisions.

At OpenAI, we are committed to the principles of broad access, collaboration, and safety for beneficial AI. As we continue to pursue our mission, we recognize the role that patents play in the technology landscape, and commit to using our patents in a way that supports innovation and our dedication to our mission.

The opening of OpenAI's pledge establishes two foundational elements: first, a statement of principles ('broad access, collaboration, and safety for beneficial AI') and second, an acknowledgment of patents' role coupled with a commitment to align patent use with these principles. This structure mirrors a practice followed in several patent pledges, where general principles typically precede specific commitments,⁹ although here the principles remain

anchored to the company's broader mission rather than creating standalone obligations.

We support the efforts of others in the use and development of AI model technology. We pledge to only use our patents defensively...

OpenAI pledges that it will not proactively assert its patents against other parties. Instead, it will only assert its patents defensively. This language implies a 'non-aggression' stance, which could make this entity more attractive as a collaborative partner. However, it does not prevent the entity from defending its IP if threatened. The operational provisions begin with a statement of support for 'the efforts of others in the use and development of AI model technology', followed by the core commitment to 'only use our patents defensively'. This formulation differs from other prominent patent pledges in both structure and substance. For instance, while IBM's open source pledge from 2005 contains an explicit commitment 'not to assert' specified patents,¹⁰ and Moderna's pledge stated it 'will not enforce our COVID-19-related patents', OpenAI's pledge is framed in terms of defensive use rather than non-enforcement. This distinction in language is significant. A commitment to 'only use patents defensively' appears to relinquish the right to initiate patent proceedings offensively while preserving the ability to use patents in response to actions by others. In practical terms, this suggests that developers working on AI model technology need not worry about unprovoked patent assertions from OpenAI. However, the pledge retains OpenAI's ability to use its patents as a shield rather than surrendering enforcement rights entirely.

The pledge's scope covers 'AI model technology'—a term that, like IBM's reference to activities related to Open Source Software, appears intentionally broad to encompass a wide range of development and implementation activities in the field of artificial intelligence models.

... so long as a party does not threaten or assert a claim, initiate a proceeding, help someone else in such activities against us, or engage in activities that harm us or our users.

The pledge's operational scope is defined through four distinct conditions that, if triggered, would release OpenAI from its commitment to defensive-only patent use: threatening or asserting a claim, initiating a proceeding, helping someone else in such activities against OpenAI and engaging in activities that harm OpenAI or its users. This sets conditions under which the entity may forgo the defensive posture and exercise its patents. If a third party initiates or threatens legal action, assists another party in a legal action against the pledging entity or asserts any form of claim (which could be interpreted broadly, potentially including cease-and-desist letters or formal notices of infringement), the pledge does not apply. These conditions serve as a protective mechanism for the pledging entity, incentivizing other parties to avoid hostile actions.

The first three conditions relate to legal or quasi-legal actions, creating a cascade of triggering events that extends from direct

⁵ JL Contreras, 'The Evolving Patent Information Landscape' (2018) CIGI 166, 1–12.

⁶ JF Ehrnsperger and F Tietze, 'Patent Pledges, Open IP, or Patent Pools? Developing Taxonomies in the Thicket of Terminologies' (2019) 14 *PLoS One* e0221411.

⁷ See JL Contreras, 'Patent Pledges as Portfolio Management Tools: Benefits, Obligations and Enforcement' in K Blind and N Thumm (eds), *A Modern Guide to Patenting: Challenges of Patenting in the 21st Century* (Edward Elgar, forthcoming, 2024); JL Contreras, 'Patent Pledge Enforcement Theories' in JL Contreras and M Jacob (eds), *Patent Pledges: Global Perspectives on Patent Law's Private Ordering Frontier* (Edward Elgar 2017) 191; BM Hill, 'Powering Intellectual Property Sharing: How to Make Tesla's Patent Pledge Effective' (2016) 24 *J Intell Prop L* 191.

⁸ Lenarczyk, Minssen and Aboy (n 1).

⁹ Other recent patent pledges follow this structural approach. OVHcloud's pledge opens with its vision of 'an open, sovereign cloud' before detailing commitments; IKEA grounds its pledge in the belief that 'everyone has the right to feel safe'; Journey Colab establishes its ethical position on psychedelic therapeutics and the Low Carbon Patent Pledge begins with its commitment to 'being a force for good through open innovation'.

See 'Patent Pledge' (OVHcloud) <<https://www.ovhcloud.com/en/about-us/patent-pledge/>> accessed 22 January 2025; 'Patent Pledge' (IKEA) <<https://www.ikea.com/global/en/patent-pledge/>> accessed 22 January 2025; 'Patent Pledge' (Journey Colab) <<https://www.journeycolab.com/patent-pledge>> accessed 22 January 2025; 'The Pledge' (Low Carbon Patent Pledge) <<https://lowcarbonpatentpledge.org/the-pledge/>> accessed 22 January 2025.

¹⁰ JF Ehrnsperger, 'Patent Pledges—Dataset' (2019) <<https://doi.org/10.17863/CAM.41462>> accessed 22 January 2025.

Table 1. Comparative analysis of Moderna and OpenAI's patent pledges

Aspect	Moderna case analysis	Application to OpenAI's pledge
Period of applicability	Defined by pandemic duration, but lack of specificity led to court interpretation using the WHO's pandemic timeline. Clear markers could reduce ambiguity.	No explicit timeframe or external marker, creating potential for indefinite applicability. Defining temporal limits could prevent reliance-based disputes.
Nature of pledge	The court treated it as 'simple consent' rather than a contract or federal law waiver due to lack of offer, acceptance or formal relinquishment of rights.	Appears as 'simple consent' based on a conditional defensive stance, with no offer or acceptance. This allows flexibility but limits enforceability as a binding commitment.
Revocability	Language allowed for updates or revocation, as demonstrated by Moderna's March 2022 update. Lacks irrevocability, enabling Moderna to adapt based on conditions.	Conditions imply revocability (eg 'harm' clause). OpenAI can likely retract the pledge if conditions are breached, similar to Moderna, creating flexibility but potential reliance concerns.
Conflict of laws	Lacked governing law clause, leading to complex jurisdictional issues. The court relied on Massachusetts and UK law but noted importance of jurisdiction choice.	No stated jurisdiction, raising cross-border interpretation issues. Specifying a governing law could streamline legal clarity, especially for international collaborators relying on the pledge.
Scope and targeted recipients	Broad pledge scope ('anyone making pandemic vaccines') led to disputes over coverage. The court advised clearer inclusion/exclusion criteria to reduce ambiguities.	Scope based on recipient behaviour (eg no hostile actions), which may be interpreted variably. Specifying intended audience or 'harmful' actions could clarify who benefits from the pledge.
Legal interpretation and enforceability	Treated as non-binding, revocable consent rather than a contractual commitment, limiting enforceability for third parties without further agreements.	Likely viewed as revocable simple consent rather than a binding contractual commitment, reducing contractual enforceability for third-party reliance but allowing OpenAI to retain control over IP commitments.
Strategic implications	Allowed Moderna to support public health without binding IP surrender, balancing flexibility with corporate interests. Flexibility limited third-party reliance.	Supports AI innovation while retaining control over proprietary rights. Conditional language (eg defensive use) promotes collaboration while guarding against potential threats.

legal action to auxiliary support of such actions. The inclusion of 'threaten' alongside actual assertion of claims suggests that the condition may be triggered before formal legal proceedings commence.

The fourth condition—"activities that harm us or our users"—stands apart in its breadth. Unlike Moderna's pledge, which was bounded by a definable temporal condition ("while the pandemic continues"), this condition introduces a qualitative assessment of harm that could encompass a range of activities beyond patent disputes. The broad wording provides flexibility, allowing the pledging entity to protect against potentially harmful uses of its technology without specifying what constitutes 'harm'. However, the ambiguity could lead to legal disputes over what actions are sufficiently 'harmful' to justify the use of patents, as courts could interpret this differently based on context.

Other companies, like Tesla and IBM, have issued patent pledges to promote industry standards and foster technology adoption.¹¹ However, this pledge is more restrictive than some others due to the added conditions that allow defensive use if any 'harm' occurs, making it a semi-defensive strategy. This could be more favourable to OpenAI but might deter some collaboration from cautious third parties.

The pledge's operation raises several interpretative considerations: the precise moment of the pledge's revocation upon a triggering event, the continuing availability of pledge-based defences up to the point of breach and the temporal reach of revocation. Additional complexity stems from the breadth and interrelation of the triggering conditions. The pledge's party-specific language

('a party') also introduces questions about whether triggering events operate individually or collectively across implementers. These operational aspects will be examined in detail through the interpretative framework established in *Moderna v Pfizer*.

III Comparative analysis of Moderna's and OpenAI's patent pledges

The *Moderna* case provides valuable insights into interpreting and structuring patent pledges. OpenAI's pledge, like Moderna's, is positioned as non-binding, conditional and potentially revocable. The open-ended timeframe, lack of specific recipients and absence of a governing law clause create flexibility but may reduce the reliability and enforceability of the pledge for third parties.

Table 1 and this section provide a brief comparative analysis of Moderna's and OpenAI's patent pledges, and show how the *Moderna* ruling can be applied to analyse the potential legal effects stemming from OpenAI's pledge.

A. Period of applicability

In *Moderna*, the court's interpretation of the pledge's duration was central, as the pledge's language ('while the pandemic continues') left ambiguity about its endpoint. The court ultimately tied the duration to the World Health Organization's official pandemic declaration, rejecting Moderna's narrower interpretation. The case demonstrates that temporal conditions should reference objectively determinable external markers to avoid disputes.

Unlike Moderna's pledge, OpenAI's pledge lacks an explicit timeframe or external markers, which could imply indefinite applicability. This absence may introduce ambiguity, leaving the duration open to interpretation and potentially enabling OpenAI to terminate the pledge without clear guidelines. To address

¹¹ Tesla, 'Patent Pledge' <<https://www.tesla.com/legal/additional-resources-intellectual-property>> accessed 29 November 2024. While Tesla's pledge was publicly framed as promoting electric vehicle technology adoption, scholars have noted it likely serves multiple strategic purposes. See, eg, JL Contreras, 'Patent Pledges' (2015) 47 *Arizona State Law Journal*, University of Utah College of Law Research Paper No 93.

this, OpenAI might consider adopting a defined period or specific criteria (eg industry or technological standards) to mitigate reliance-based disputes, especially for long-term adopters.

B. Nature of the pledge

The court in *Moderna* interpreted the pledge as 'simple consent' rather than a unilateral contract or federal law waiver. The absence of offer, acceptance and consideration pointed to a non-binding, temporary stance rather than a contract. Moderna's language also avoided federal law implications, which would have required an explicit waiver of rights.

OpenAI's language suggests a similar 'simple consent' approach. By pledging only to use its patents defensively unless certain conditions are breached, OpenAI's stance resembles Moderna's non-binding, temporary position rather than a contractual commitment. Like Moderna, OpenAI does not seem to aim to create a formal contract or waiver, which maintains flexibility but potentially limits enforceability for third parties relying on this pledge for business planning or innovation.

C. Revocability

Moderna's language, coupled with forward-looking statements, allowing for updates or revocation, was confirmed by its March 2022 update, which the court deemed a valid notification of change. This flexibility highlighted that, unless explicitly irrevocable, patent pledges can be withdrawn based on changing conditions, potentially complicating reliance by third parties.

OpenAI's pledge, contingent upon conditions related to defensive use, appears similarly revocable. By tying the pledge to certain behaviours ('as long as a party does not threaten, assert a claim ... or engage in activities that harm us or our users'), OpenAI implies that any breach of these conditions could justify revocation. This flexible stance is strategically useful for OpenAI but could deter reliance by other entities that prefer long-term stability.

D. Conflict of laws

In *Moderna*, the absence of a governing law clause raised complex jurisdictional questions. The court analysed the pledge under Massachusetts and UK law, depending on whether it was interpreted as a unilateral contract or a federal law waiver. This underscores the importance of a specified governing law, as it determines interpretation and enforceability, especially across borders.

OpenAI's pledge, which lacks a stated jurisdiction, could face similar interpretative challenges. As a US-based company, its pledge might be governed by US federal or state law if interpreted in a US court, but international disputes could result in inconsistent application. Specifying a governing law could streamline legal interpretation and provide greater predictability for international adopters.

E. Scope and targeted recipients

The Moderna pledge, aimed at 'anyone making vaccines intended to combat the pandemic', was broad and created inclusion ambiguities. Moderna's language did not specify intended beneficiaries, leading to disputes over whether certain entities were covered. The court ultimately interpreted the pledge as applying to a wide audience but noted that clearer language would have minimized disputes.

OpenAI's pledge avoids directly naming intended recipients, instead relying on behaviours to limit who may benefit from the pledge (eg those who do not initiate claims or cause harm). This broad, behaviour-based scope might create ambiguity, as it

relies on subjective criteria like 'harm', which could be interpreted differently in various jurisdictions. Clarifying the intended audience or specifying actions that qualify as harmful could reduce uncertainty for parties seeking to rely on OpenAI's pledge.

F. Legal interpretation and enforceability

Moderna's pledge was ultimately treated as a temporary, non-binding consent rather than an enforceable contract. The case emphasized that, without offer and acceptance, a patent pledge may not create binding obligations.

Similarly, OpenAI's pledge lacks contractual elements and appears intended as simple consent rather than a binding commitment. Without offer, acceptance or explicit consideration, OpenAI's pledge likely limits *contractual* enforceability for third parties who may rely on it. This approach may be suitable for OpenAI, as it avoids long-term commitments, but it could deter third parties from relying on the pledge for substantial investment or development based on OpenAI's IP.

IV Post-Moderna analysis of OpenAI's pledge

A. Nature of the pledge: consent, unilateral contract or federal law waiver

Applying the High Court's framework in *Moderna*, OpenAI's pledge operates as 'simple consent' to what would otherwise be infringing acts.¹² The commitment to 'only use our patents defensively' manifests present intent regarding patent enforcement, meeting the threshold established in *Moderna* for conferring legally meaningful consent.

While OpenAI's pledge employs defensive language rather than Moderna's explicit non-enforcement commitment, this distinction does not alter its fundamental character as consent under the High Court's analysis. The essence of both statements is the relinquishment of proactive patent enforcement rights against specified activities. As Richards J explained, when a patentee states it will not enforce patents in specified circumstances, 'the fair interpretation of such a statement is that the other may exploit the patent without consequence since the monopoly rights that are the essence of the patent are not being asserted'¹³.

The pledge incorporates defensive use rather than non-enforcement, operates subject to behavioural conditions and applies on a party-specific basis. Following *Moderna*, the pledge is unlikely to constitute a unilateral contract, as the language contains no objective manifestation of intent to form contractual relations. The inclusion of behavioural conditions and the party-specific nature of the pledge do not elevate it to contractual status, as the High Court emphasized that even conditional pledges can operate as simple consent.

Similarly, the pledge does not amount to a federal law waiver under the *Moderna* framework. The High Court distinguished between temporary forbearance to sue (which can constitute simple consent) and the permanent relinquishment of rights required for a federal law waiver. OpenAI's pledge, with its conditional nature and potential for modification, falls squarely within the former category.

Like Moderna's pledge, this represents a present position capable of modification rather than an irrevocable commitment. The consent operates without the need for consideration or reliance by

¹² *Moderna v Pfizer*, paras 183–87 (establishing that public statements by patentees not to enforce patents can constitute consent to what would otherwise be infringing acts under s 60 of the Patents Act). See also Lenarczyk, Minssen and Aboy (n 1).

¹³ *ibid*, para 183.

implementers, creating potential defences against infringement claims while the pledge remains in effect. This characterization as simple consent under the *Moderna* framework informs the analysis of the pledge's conditions, revocation mechanisms and jurisdictional implications in subsequent sections.

The High Court's analysis suggests that such consent, while legally meaningful, remains subject to the pledgor's ability to modify or revoke it. This aligns with the flexible nature of patent pledges as instruments that can accommodate both the pledgor's strategic interests and the public benefit of reduced patent assertion, without requiring the formality or permanence of contractual commitments.

B. Scope and period of applicability

OpenAI's pledge creates a simple consent defence against patent infringement claims while simultaneously establishing boundaries for that consent through its behavioural conditions. The consent operates to prevent proactive patent enforcement by OpenAI against parties developing AI model technology but ceases upon specific triggering events.

The primary conditions governing revocation fall into distinct categories. The first three conditions—threatening claims, initiating proceedings or helping others in such activities—create bright-line triggers based on identifiable legal or quasi-legal actions. Following *Moderna*'s approach, these conditions would operate from the moment such actions commence.

The 'harm' condition introduces greater complexity and requires factual assessment. The High Court's analysis in *Moderna* suggests such conditions remain enforceable despite their interpretative challenges, provided they establish objectively determinable boundaries.¹⁴

The conditions create an asymmetric structure of protection: while 'harm' may require factual determination and could operate prospectively from the point of established breach, the legal action conditions (threatening or asserting claims) appear to make the consent void from inception. This aligns with the pledge's defensive nature, creating stronger consequences for direct legal challenges than for general harmful conduct.

C. Revocability

As a form of simple consent under the *Moderna* framework, OpenAI's pledge embodies a dual revocation structure: a unilateral right to withdraw consent and automatic cessation triggered by specified behavioural conditions. This binary structure presents novel temporal and operational complexities that merit examination.

The automatic revocation mechanism operates distinctly from a unilateral revocation. When a party breaches the pledge's conditions—whether by initiating proceedings against OpenAI or engaging in activities that constitute 'harm'—the pledge's protection for that party may be voided from inception, as if the consent had never been granted. This differs fundamentally from the scenario where OpenAI exercises its right to unilaterally revoke the pledge. Under *Moderna*'s analysis, such unilateral revocation would operate prospectively only—activities conducted before revocation would retain their defence against infringement claims.

While the legal actions (threats, claims and proceedings) provide clear triggering events, the 'harm' condition introduces interpretative complexity. The fundamental question is not temporal

but substantive: what constitutes 'harm' to OpenAI or its users? Market competition illustrates this challenge. Could successful competition in the AI market constitute 'harm' to OpenAI? A competing product gaining market share might impact OpenAI's commercial interests, but whether this amounts to 'harm' under the pledge is unclear. Depending on interpretation, competitive impact could theoretically be deemed harmless or, conversely, be viewed as harm from the moment the competing product entered the market.

The temporal implications extend to damages calculations and vary between the two revocation mechanisms. For unilateral revocation, the calculation is straightforward—damages would only be available for post-revocation activities. However, for breaches of conditions, if the pledge is void *ab initio*, damages might theoretically be available for all infringing activities, regardless of when the breach was discovered.

The party-specific nature of the conditions adds another layer of complexity. Under *Moderna*'s framework, consent operates on a targeted rather than universal basis. Thus, OpenAI's pledge likely creates individual relationships with each implementer, allowing revocation to operate independently for different parties. This granular operation aligns with the pledge's defensive structure—consent remains valid for parties not triggering conditions, even if others have breached them.

Fundamentally, these temporal and operational questions reflect the tension between the pledge's defensive nature and its behavioural conditions. The framework established in *Moderna* provides guidance but requires adaptation to address the distinct challenges posed by conditions that may lack clear definitional boundaries.

D. Conflict of laws

The conflict of laws analysis for OpenAI's pledge mirrors the complexity encountered in *Moderna*.¹⁵ Like *Moderna*'s pledge, OpenAI's statement lacks an explicit choice of law clause and operates globally across multiple jurisdictions. The *Moderna* judgment clarifies that, when a pledge functions as simple consent, its effect as a defence to patent infringement will be determined by the national patent law of each jurisdiction where the patent rights exist. This territorial approach reflects the fundamental nature of patent rights as nationally granted monopolies.

E. Legal interpretation and enforceability

OpenAI's statement constitutes a patent pledge by meeting the core definitional elements—it is a voluntary public commitment to limit patent enforcement through defensive-only use. Operating within the interpretative principles established in *Moderna*, the pledge's effectiveness does not depend on precise drafting or exhaustive definitions.¹⁶ Rather, the key question is whether it manifests a clear present intention regarding patent enforcement. OpenAI's commitment to 'only use our patents defensively' satisfies this threshold, creating a defence against infringement claims while the pledge remains operational. This defence persists until either OpenAI explicitly revokes the pledge or an implementer triggers one of the behavioural conditions.

¹⁵ *ibid*, paras 66–75 (examining application of Rome I and Rome II to determine governing law for patent pledges lacking explicit choice of law provisions). See also Lenarczyk, Minssen and Aboy (n 1).

¹⁶ *ibid*, paras 33–34 and 57 (establishing that patent pledges are interpreted objectively based on how a reasonable reader would understand them, while recognizing they can serve as initial positions for negotiations). See also Lenarczyk, Minssen and Aboy (n 1).

¹⁴ *ibid*, paras 33–35 (discussing interpretation of conditions in patent pledges through an 'ordinary objective meaning/reasonable reader' standard and confirming their enforceability). See also Lenarczyk, Minssen and Aboy (n 1).

The pledge's enforceability operates asymmetrically across its conditions. The legal action triggers (threatening claims and initiating proceedings) create objectively determinable boundaries that, when breached, may void the consent from inception. The 'harm' condition, while broader, remains enforceable under Moderna's principles but requires factual determination of the occurrence of harmful activities to establish if consent ceased.

The pledge's defensive formulation and behavioural conditions do not diminish its legal effect while operational. Following *Moderna*, implementers can rely on the pledge as a defence against infringement claims for activities undertaken before valid revocation. However, this protection remains contingent on adherence to the pledge's conditions—any breach not only terminates future protection but may retrospectively make the consent void from inception.

The result is an enforceable commitment that provides implementers' protection against unprovoked patent assertions while maintaining OpenAI's ability to respond to adverse actions.

V. Discussion and practice points

The analysis of OpenAI's patent pledge, viewed through the interpretative lens of the *Moderna v Pfizer/BioNTech* decision, highlights the legal and strategic complexities inherent in modern patent pledges.

OpenAI's approach—based on defensive use subject to behavioural conditions—reflects an evolution in the formulation of these commitments. Unlike Moderna's temporal condition tied to the duration of a public health emergency, OpenAI's reliance on behavioural triggers introduces interpretative challenges, particularly concerning the scope of 'harm' and the exact moment of revocation. While this flexibility aligns with OpenAI's strategic goals, it raises questions for third parties about the long-term reliability of such a pledge.

From a practical standpoint, the absence of explicit temporal limits, defined beneficiaries and a governing law clause could create significant uncertainties in both domestic and cross-border applications of OpenAI's pledge. The *Moderna* judgment underscored the necessity of clear and precise language to enhance enforceability and prevent disputes, a lesson that applies directly to OpenAI's formulation. These ambiguities could discourage reliance by smaller innovators and collaborators, reducing the overall utility of the pledge as a tool for fostering open innovation in the AI sector.

To strengthen the effectiveness of patent pledges and ensure they achieve their dual objectives of promoting innovation and protecting proprietary interests, several measures

are recommended. OpenAI and similar pledgors should consider adopting explicit temporal conditions, narrowing the scope of triggering behaviours and clearly specifying governing law. Such adjustments would reduce interpretative ambiguities and provide greater assurances to collaborators, creating a more stable environment for technological development. By addressing these elements, patent pledges can evolve into robust instruments that balance flexibility with legal certainty, advancing innovation in fast-moving fields such as artificial intelligence.

That said, strategic ambiguity in patent pledges, such as OpenAI's, can also offer significant advantages, particularly in fields characterized by rapid technological change and intense competition. By avoiding overly rigid formulations, OpenAI retains the flexibility to adapt its pledge to unforeseen circumstances, protecting its proprietary interests while maintaining a public commitment to fostering innovation. This flexibility allows OpenAI to respond dynamically to evolving market and legal conditions, a critical consideration in the context of artificial intelligence, where the pace of development often outstrips regulatory and contractual frameworks.

Moreover, ambiguity in the pledge's terms can encourage implementers to engage in good-faith collaborations, knowing that compliance with behavioural conditions ensures protection under the pledge. This conditional approach promotes trust while providing OpenAI with a safeguard against opportunistic actions that might undermine its mission or users' interests. Strategic ambiguity, when employed thoughtfully, can thus create a balanced ecosystem that incentivizes responsible innovation without overburdening the pledgor with irrevocable commitments.

Ultimately, the success of patent pledges like OpenAI's lies in achieving an optimal balance between flexibility and predictability. While greater clarity can enhance legal certainty, strategic ambiguity, when transparently communicated, can foster collaboration while preserving the pledgor's ability to protect its core values and respond to challenges. OpenAI's pledge, with its nuanced approach, illustrates the potential of patent pledges to operate as dynamic tools that adapt to the complexities of modern technological ecosystems.

Funding

The research was supported, in part, by a Novo Nordisk Foundation Grant for a scientifically independent International Collaborative Bioscience Innovation & Law Programme (Inter-CeBIL programme, grant no. NNF23SA0087056).