

**AI & ANTITRUST: RECONCILING TENSIONS BETWEEN
COMPETITION LAW AND COOPERATIVE AI DEVELOPMENT**

Shin-Shin Hua¹ & Haydn Belfield²

Cooperation between companies developing artificial intelligence (AI) can help them create AI systems that are safe, secure, and with broadly shared benefits. Researchers have proposed a range of cooperation strategies, ranging from redistributing “windfall” profits to assistance to address the harmful dynamics of a competitive race for technological superiority.

A critical tension arises, however, between cooperation and the goal of competition law, which is to protect the very process of competition between rival companies. Whilst these potential conflicts are significant, they are currently underexplored in the literature. This paper examines the relationship between proposed forms of AI cooperation and competition law, focusing on the competition law of the European Union (EU).

EU competition law governs the behavior of the world’s largest AI companies, though many are based abroad, especially in

¹ Research Affiliate, Centre for the Study of Existential Risk, University of Cambridge; antitrust attorney at British Telecom Group and formerly at Cleary Gottlieb Steen & Hamilton, London. Shin-Shin has an Advanced Masters in Public International Law from Leiden University (Cum Laude) and a degree in Law from Peterhouse, University of Cambridge.

² Research Associate and Academic Project Manager at the Centre for the Study of Existential Risk, and Associate Fellow of the Leverhulme Centre for the Future of Intelligence, both at the University of Cambridge. Haydn is a former Senior Parliamentary Researcher to a British Shadow Cabinet Minister and Policy Associate to the University of Oxford’s Global Priorities Project. He has a degree in Philosophy, Politics and Economics from Oriel College, University of Oxford.

the US. Its jurisdiction can extend to any foreign company that is active in the EU. Scrutiny of US “Big Tech” is also an area of strategic focus for the European Commission (EC).

This paper seeks to reconcile the cooperative AI development and competition law. It examines fourteen forms of AI cooperation, both those that are applicable today and longer-term strategies that will apply when AI development is more advanced. Where we identify potential tensions with EU competition law, we suggest mitigation steps. Our aim is to ensure the long-term sustainability of these important safeguards to the responsible and beneficial development of AI.

Table of Contents

Introduction.....418

I. Why is Cooperative AI Development Desirable?428

II. Background: The Prohibition on Anti-Competitive Agreements
Under Article 101433

 A. Article 101(1) Prohibition.....434

 B. Article 101(3) Exemption.....439

 C. Investigation and Enforcement Powers.....444

III. Legal Analysis of Proposed Forms of Cooperation: Long-Term
Strategies454

 A. OpenAI Assist Clause455

 B. The Windfall Clause.....482

IV. Legal Analysis of Proposed Forms of Cooperation: General
Trust Building.....489

 A. Agreements Between Competitors.....490

 B. Information Exchange492

 C. Standard-Setting.....516

Conclusion537

INTRODUCTION

This paper surveys types of cooperation between AI companies and proposes ways to structure cooperation to comply with competition law.

AI governance³ is concerned that without cooperation between AI companies, AI development could be less safe and beneficial. Researchers have therefore proposed many cooperation strategies, from the Assist Clause to information exchange and standard setting. However, competition law could be a barrier to these strategies. This is because, at its heart, competition law seeks to promote competition and prevent cooperation that harms consumers. Nevertheless, if structured carefully, these cooperation strategies should not raise serious competition law concerns, and so achieve their objectives of cooperation and risk reduction.

In this paper, we define “AI” as a digital system that is capable of performing tasks commonly thought to require intelligence, with these tasks typically learned via data and/or experience.⁴ “AI system” refers to a software process (with the

³ The study or practice of local and global governance systems—including norms, policies, laws, processes, and institutions—that govern or should govern AI research, development, deployment, and use. Sean S. Ó hÉigartaigh, Jess Whittlestone, Yang Liu et al. *Overcoming Barriers to Cross-cultural Cooperation in AI Ethics and Governance*, PHILOS. TECH. (May 15, 2020), <https://link.springer.com/content/pdf/10.1007/s13347-020-00402-x.pdf>.

⁴ Miles Brundage, Shahar Avin, Jasmine Wang, Haydn Belfield, Gretchen Krueger et al., *Toward Trustworthy AI Development: Mechanisms for Supporting Verifiable Claims*, arXiv preprint arXiv:2004.07213 (Apr. 20, 2020) <https://arxiv.org/pdf/2004.07213.pdf>, at 4.

characteristics of AI mentioned above), running on physical hardware, under the direction of humans operating in some institutional context.⁵

Companies play an outsized role in AI research and development compared to academic or government groups. Cooperation between competing companies is likely necessary in order for them to develop and deploy AI systems *responsibly*—that is, safely, securely and with broadly distributed benefits.⁶ On the other hand, lack of cooperation could be *irresponsible*: unsafe, insecure, and not socially beneficial. AI governance is especially concerned with the prospect of AI research taking on the dynamics of a competitive race for technological superiority.⁷ These considerations have led to cooperation between AI companies becoming a major focus of AI governance, at both theoretical and practical levels.

⁵ *Id.* at 62

⁶ Amanda Askell, Miles Brundage & Gillian Hadfield, *The Role of Cooperation in Responsible AI Development*, COMPUTERS AND SOCIETY, arXiv preprint arXiv:1907.04534 (2019). Allan Dafoe, Edward Hughes, Yoram Bachrach, Tantum Collins, Kevin R. McKee, Joel Z. Leibo, Kate Larson & Thore Graepel, *Open Problems in Cooperative AI*, arXiv preprint arXiv:2012.08630 (15 Dec 2020).

⁷ See generally *id.*; Nick Bostrom, Allan Dafoe & Carrick Flynn, *Policy Desiderata in the Development of Machine Superintelligence*, FUTURE OF HUMANITY INSTITUTE OXFORD (FHI) WORKING PAPER (2018); Stuart Armstrong, Nick Bostrom & Carl Shulman, *Racing to the Precipice: a Model of Artificial Intelligence Development*, 31 AI & SOCIETY 201 (2016); Stephen Cave & Sean S. Ó hÉigeartaigh, *An AI Race for Strategic Advantage: Rhetoric and Risks*, 2018 PROCEEDINGS OF THE 2018 AAAI/ACM CONFERENCE ON AI, ETHICS AND SOCIETY 36; Haydn Belfield, *Activism in the AI Community: Analysing Recent Achievements and Future Prospects*, 2020 PROCEEDINGS OF THE 2020 AAAI/ACM CONFERENCE ON AI, ETHICS AND SOCIETY 15; Ó hÉigeartaigh, Whittlestone, Liu *et al.*, *supra* note 3.

However, competition law could raise barriers to these proposed strategies. There is an inherent tension between cooperation and the goals of competition law, which at its core is meant to protect the very processes of rivalry between companies. Strategies that seek cooperation rather than competition between companies can therefore raise anti-competitive concerns. In particular, these strategies may breach the EU competition rule that prohibits agreements between companies that restrict or distort competition between them: Article 101(1) of the Treaty on the Functioning of the European Union (TFEU). This rule is the counterpart of Section 1 of the Sherman Act in the US, which declares that “[e]very contract, combination in the form of trust or otherwise, or conspiracy, in restraint of trade or commerce” is illegal.⁸

The scope of this note focuses on the EU competition regime, with some limited references to the UK regime where relevant. While many leading AI companies are based in the US or China, the practical reach of EU competition law is broad and any foreign company that has activities or customers in the EU, or any other impact on EU trade, can fall within the jurisdiction of EU competition law. EU competition law is powerful and influential. Breaches can lead to fines of up to 10% of worldwide turnover, for

⁸ 15 U.S.C. § 1.

example, as well as derivative private litigation claims for sums that can dwarf the original competition law fine.⁹ Even if investigations and fines were only future possibilities, the legal uncertainty could dissuade companies from participating in types of cooperation. Cooperation that raises competition concerns in an EU context is likely to raise similar concerns in the US or beyond. Recent European cases in the technology sector have influenced similar investigations across the world.¹⁰ Therefore, it is important that we structure cooperation in ways that do not raise EU competition law concerns. We suggest how to do so below.

The purpose of this paper is not to survey the incentives of AI companies to cooperate or the likelihood that they may do so. Rather, it is to examine initiatives that the AI research community is already discussing and AI labs are already employing. We think cooperative AI development is feasible despite fierce market rivalry

⁹ In the UK, for example, the current live claim for *Merricks v Mastercard* seeks £14 billion in damages, brought on behalf of 46 million consumers (*Mastercard Incorporated and others (Appellants) v Walter Hugh Merricks CBE (Respondent)* [2020] UKSC 51). The original competition law decision (Case COMP/34.579-*Mastercard*, EC decision of 19 December 2009), upheld on appeal, did not impose a fine but the defendant was obliged to bring the infringement to an end.

¹⁰ For example, in 2005, South Korea closely followed the EU's ruling in 2004 against Microsoft relating to its bundling of Windows Media Player with its operating system (Case COMP/C-3/37.792, *Microsoft*, Commission decision of 24 May 2004, 2007 O.J. L32/23). The EU's multiple investigations into Google since 2010 (Case AT.39740, *Google Shopping*, Commission decision of 27 June 2017; Case AT.40099, *Google Android*, Commission decision of 12 February 2018; Case AT.40411, *Google AdSense*, Commission decision of 20 March 2019) have also led to similar investigations in Russia's in 2015 regarding Google Android, Brazil in 2013 for Google Search and later for Android, Turkey in 2015 for Google Android, and South Korea in 2016 for Google Search and Android. For further discussion, see ANU BRADFORD, *THE BRUSSELS EFFECT: HOW THE EUROPEAN UNION RULES THE WORLD*, at 122-27.

between the world's largest AI companies today for two main reasons.

First, governments, civil society and public opinion currently focus on the need for safe, beneficial and responsible AI. See, for example, the European Commission's proposed Artificial Intelligence Act, that is currently passing through the EU legislative process and seeks to build trust in AI systems and ensure they are used consistently with fundamental rights and European values.¹¹ Responsible AI development is an essential competitive parameter for AI companies. Conversely, AI companies that do not meet expected standards can experience severe commercial and reputational harm, as most vividly demonstrated by the Cambridge Analytica scandal.¹²

Second, cooperation between AI companies need not be a zero-sum game. It can be beneficial for innovation, interoperability, consumer trust and adoption to cooperate. Incident sharing and establishing industry safety benchmarks can be mutually beneficial to all players, for example. We also see robust industry cooperation today between companies and in bodies like the Partnership on AI

¹¹ Proposal for a Regulation of the Euro. Parliament and of the Council Laying Down Harmonised Rules on Artificial Intelligence (Artificial Intelligence Act) and Amending Certain Union Legislative Acts, Com/2021/206 final.

¹² See, e.g., Nicholas Confessore, *Cambridge Analytica and Facebook: The Scandal and the Fallout So Far*, N.Y. TIMES (April 4, 2018), <https://www.nytimes.com/2018/04/04/us/politics/cambridge-analytica-scandal-fallout.html>.

(PAI), which includes a wide variety of the world's largest AI and technology companies.¹³

It is true, though, that some of the longer-term cooperation strategies will require companies to sacrifice a large amount of profits (Windfall Clause) or effectively shut down or fundamentally change their business (Assist Clause).¹⁴ These sacrifices can be justified by the extreme risks associated with advanced AI systems, as we will go on to discuss. Indeed, several leading AI companies have committed to the goal of ensuring the safe and beneficial development of AI, such as DeepMind (part of Google) and OpenAI.

Scope and Contribution

The tensions between AI strategies¹⁵ and competition law are potentially significant, yet currently underexplored in AI governance. This Article builds on work at this intersection.¹⁶ We

¹³ PAI is a non-profit organization that brings together 100 'Partners' including the leading Western technology companies, NGOs such as Amnesty, and academic groups. It shares best practice, conducts research and XX. It was founded in 201X and is based in San Francisco.

¹⁴ For further discussion of the distinction between near- and longer-term, *see infra* Section I (Why is Cooperative AI Development Desirable?).

¹⁵ AI strategy is the study or practice of high-level plans for how specific private- or public-sector actors can achieve their AI-related goals. *See* Ó hÉigearthaigh, Whittlestone, Liu *et al.*, *supra* note 3.

¹⁶ *See generally* Miles Brundage, Shahar Avin, Jack Clark *et al.*, *The Malicious Use of Artificial Intelligence: Forecasting, prevention, and mitigation*, arXiv:1802.07228 (Feb. 18, 2018) <https://arxiv.org/pdf/1802.07228.pdf>; Askill, Brundage & Hadfield, *supra* note 6; Haydn Belfield, *From Tech Giants to a Tech Colossus: Antitrust Objections to the Windfall Clause* (Mar. 15, 2019) (unpublished manuscript) (on file with author); Brundage, Avin, Wang *et al.*, *supra* note 4; Dakota Foster & Zachary Arnold, *Antitrust and AI: How Breaking Up Big Tech Could Affect the Pentagon's Access to AI*, CENTER FOR SECURITY AND EMERGING TECHNOLOGY (May 2020); Cullen O'Keefe,

also draw on the academic literature¹⁷ and relevant jurisprudence¹⁸ of relevant areas of competition law, including its future direction of travel in regulating emerging technology and AI.¹⁹

Based on our legal analysis, the objective of this paper is to advise on whether forms of AI cooperation raise competition law risks,²⁰ and if so how these competition law risks can be practically mitigated. We assume that we want both to uphold competition law, and for these forms of AI cooperation to be lawful and to succeed. This paper does not comment on whether the competition law regime hinders or furthers the aims of AI governance more broadly. Neither does it attempt to compare the potential benefits to AI governance from competition law against the costs.

How Will National Security Considerations Affect Antitrust Decisions in AI? An Examination of Historical Precedents, FHI WORKING PAPER (forthcoming); Sophie-Charlotte Fischer *et al.*, *Levers of Influence of the USG on AI Development*, FHI WORKING PAPER (forthcoming).

¹⁷ See *e.g.*, Lina Khan, *Amazon's Antitrust Paradox*, 126 Yale L.J. 710; Matt Stoller, *Goliath: The 100-Year War Between Monopoly Power and Democracy* (2019).

¹⁸ See *e.g.*, Case AT.40099, *Google Android*, *supra* note 10; DIGITAL COMPETITION EXPERT PANEL, UNLOCKING DIGITAL COMPETITION- REPORT OF THE DIGITAL COMPETITION EXPERT PANEL, HM TREASURY (2019), https://assets.publishing.service.gov.uk/government/uploads/system/uploads/attachment_data/file/785547/unlocking_digital_competition_furman_review_web.pdf; Elizabeth Warren, *Here's how we can Break Up Big Tech*, Medium (Mar. 8, 2019), <https://medium.com/@teamwarren/heres-how-we-can-break-up-big-tech-9ad9e0da324c>.

¹⁹ The authors combine on the one hand, nine years' experience representing clients such as NVIDIA, Google, Samsung and Sony in competition and technology cases before the EC and other global regulators; and on the other seven years' experience developing and researching policy, especially around catastrophic risk and advanced AI.

²⁰ Any reference to "risk" will refer to risk of breaching competition law, unless otherwise indicated.

Our Article suggests strategies that further the aims of both competition law and socially beneficial cooperation, seeking to reconcile the two. Our suggestions on how to reduce the risk of breaching Article 101(1)—that is, by bringing an agreement outside of the prohibition altogether, or relying on the exemption for countervailing efficiencies under (Article 101(3)) —are legitimate and lawful mitigation strategies. Our focus is on Article 101(1). However, EU competition law is broader than just this prohibition and includes rules on abuse of dominance, merger control and state aid. It follows that even if a form of cooperation complies with Article 101(1), one must still check for compliance with competition law more broadly.

Strategies for Cooperation

We examine fourteen specific proposed cooperation strategies to encourage safe and responsible AI development that are under consideration within AI governance. We group these into two broad categories: longer-term cooperation strategies intended to apply when AI development is more advanced; and nearer-term, more general cooperation strategies. For each cooperation strategy, we explain what it is, why it furthers AI governance aims, what competition concerns it might raise, and how to address those concerns.

In the coming decades, if AI continues to progress rapidly as a field, “transformative AI” may be developed and deployed. In this

Article, we use the term “transformative AI” (TAI) to refer to advanced AI systems that cumulatively could be as transformative of the international economy and society as the industrial revolution.²¹ AI has many possible uses across economy or society—it is a general-purpose technology like the steam engine or electricity, and may be as transformative. Current AI systems are still mostly “narrow,” or created specifically for a unique or targeted task or class of tasks such as voice recognition, fraud detection or vehicle automation. However, the field is working towards artificial *general* intelligence (AGI), defined by leading AI company OpenAI as “highly autonomous systems that outperform humans at most economically valuable work.”²² One can view AGI as a subset of TAI. For the purposes of our analysis, we will refer to TAI instead of AGI, as it is a broader concept that covers a wider set of AI systems.

Part I explains why cooperation is desirable for AI development, and distinguishes between forms of cooperation that are relevant in the nearer term and in longer term. Part II provides a brief background on EU competition law, including the prohibition on anti-competitive agreements, some key concepts, its investigation and enforcement powers, and an assessment of

²¹ Ross Gruetzemacher & Jess Whittlestone, *The Transformative Potential of Artificial Intelligence*, COMPUTERS AND SOCIETY, arXiv preprint arXiv:1912.00747 (2020).

²² *OpenAI Charter*, OPENAI (9 April 2018), <https://openai.com/charter/>.

relevant future developments. As it is introductory, it may be most relevant to those unfamiliar with EU competition law.

The subsequent legal analysis examines the compatibility of the fourteen forms of AI cooperation outlined above with competition law, and suggests ways to mitigate any competition law risks. Part III analyzes the Assist Clause and the Windfall Clause. These are two strategies designed to address longer-term risks. Although companies will ideally agree to be bound by these agreements today, they will not implement them until a future point in time.

Part IV examines nearer-term, general trust-building cooperation including agreements, information exchange and standard setting. Trust-building strategies include agreements between competitors, mutual monitoring and information exchanges, standard setting to proliferate best practices around an industry, and mergers.²³ To promote cooperation, companies should ideally implement these trust-building strategies now, and sustain them over the long-term trajectory of TAI development. In Part V, we conclude.

I. WHY IS COOPERATIVE AI DEVELOPMENT DESIRABLE?

Cooperation is likely necessary in order for competing artificial intelligence (AI) companies to develop AI systems *responsibly*—that is, safely, securely and with broadly distributed benefits.²⁴ We will briefly explain the types of AI risks that arise in both the nearer and longer term, and how cooperation can help to reduce them.

As discussed above, AI R&D is developing rapidly, with dramatic breakthroughs in language models, robotics, drone swarms and basic scientific research, such as protein-folding. Such progress may lead to better goods and services at lower prices, as well as significant leaps forward in science and medicine.²⁵ A leading subfield within AI is “machine learning.” Machine learning systems can iteratively improve their performance over time through experience by “learning,” that is, through identifying patterns or relationships in data with data analysis.²⁶ Many of the benefits and risks arising from AI systems come from their unique ability to “learn” in this way.

²⁴ Askill, Brundage & Hadfield, *supra* note 6. Other fields in which cooperation has proved useful include biotechnology, computer security, etc. *Publication Norms for Responsible AI: Ongoing Initiative*, PARTNERSHIP ON AI, <https://www.partnershiponai.org/case-study/publication-norms/> (last accessed Aug. 25, 2020).

²⁵ Allan Dafoe, *AI Governance: A Research Agenda*, CTR. FOR GOVERNANCE OF AI (Aug. 27, 2018), at 10.

²⁶ Shin-Shin Hua, *Machine Learning Weapons and International Humanitarian Law: Rethinking Meaningful Human Control*, 51(1) GEORGETOWN J. INTERNAT'L L. 117, at 124.

Several near-term risks from AI are relevant to our analysis. For example, AI can generate non-transparent or unexplainable outcomes. Many machine learning models generate their results in a way that is beyond the interpretive capabilities of human-scale reasoning. For example, in a model with billions of parameters it can be difficult to know which set of learnt patterns and “weights” on particular “neurons” has led to a particular outcome. In these cases, the rationale of algorithmically-produced outcomes that directly affect decision subjects remains opaque to those subjects. This opacity of the model may be deeply problematic, for example, where the processed data could harbor traces of discrimination, bias, inequity, or unfairness. This is particularly the case when people use the model to inform high consequence decisions, such as those dealing with education, credit and loans, welfare benefits, or criminal sentencing.²⁷

In some cases, AI systems may also lead to unreliable, unsafe, or poor quality outcomes. AI safety issues may arise from a variety of factors, such as irresponsible data management, poor design and production processes, and inadequate deployment practices. These outcomes can directly harm individuals and undermine public trust in potentially societally beneficial AI

²⁷ *Understanding Artificial Intelligence Ethics and Safety*, ALAN TURING INSTITUTE, at 4-5.

technologies.²⁸ In addition to the general categories of AI harm outlined above, other forms of harm from AI include amplification of bias,²⁹ loss of privacy³⁰, propagation of disinformation,³¹ and increased unemployment.³²

There is considerable disagreement over how far and how quickly the field of AI will progress, but experts believe that the capabilities of AI systems may outstrip those of humans across a range of previously hard-to-automate tasks in the coming decades.³³ As discussed above, this future form of AI- so-called TAI- could radically transform the distribution of welfare, wealth, or power, to an extent comparable to the Industrial Revolution. Potential positive changes include a significant increase in wealth, health, and well-being. TAI could also enable new forms of effective democratic decision-making and accountability, empowering democracy.

²⁸ *Id.*

²⁹ Karen Hao. *This is how AI bias really happens-and why it's so hard to fix*, MIT TECH. REV. 2019, <https://www.technologyreview.com/s/612876/this-is-how-ai-bias-really-happens-and-why-its-so-hard-to-fix/>.

³⁰ *Artificial intelligence advances threaten privacy of health data: Study finds current laws and regulations do not safeguard individuals' confidential health information*, SCIENCEDAILY (2019), <https://www.sciencedaily.com/releases/2019/01/190103152906.htm>.

³¹ Irene Solaiman et al., *Release Strategies and the Social Impacts of Language Models*, ARXIV (Aug. 2019), <http://arxiv.org/abs/1908.09203>.

³² Carl Benedikt Frey and Michael A Osborne, *The Future of Employment: How Susceptible Are Jobs to Computerisation?*, TECH. REP. THE OXFORD MARTIN PROGRAMME ON TECHNOLOGY AND EMPLOYMENT 72 (2013), <https://www.oxfordmartin.ox.ac.uk/publications/the-future-of-employment/>.

³³ Katja Grace, John Salvatier, Allan Dafoe, Baobao Zhang, and Owain Evans, *When Will AI Exceed Human Performance? Evidence from AI Experts*, ARXIV:1705.08807 [Cs] (May 24, 2017), <http://arxiv.org/abs/1705.08807>.

However, TAI could also lead to negative transformative effects. It may radically increase economic inequality or provide new tools of state repression and control. In international security, TAI could significantly alter key strategic parameters, such as the security of nuclear retaliation and the offense-defense balance. Contemporary AI systems cause safety accidents. Safety problems with advanced TAI could be much more catastrophic.

So-called “harmful race dynamics” can also arise in the development of TAI. The first groups to develop and deploy advanced TAI could gain extreme power and wealth. These rewards could drive a fierce technology race, where actors are strongly incentivized to trade-off against other values (like safety, transparency, accountability, and democracy) in order to increase the probability of gaining a competitive advantage in the race. This could catastrophically compromise responsible AI development, especially concerns around AI safety.³⁴

The complex and potentially catastrophic risks posed by AI have led to cooperation between AI companies becoming a major focus of AI governance. This paper examines fourteen concrete proposals for cooperation that encourage responsible AI development in order to mitigate both near-term and longer-term AI risks.

³⁴ Dafoe, *supra* note 25, 42-43. See also Ó hÉigeartaigh, Whittlestone, Liu *et al.*, *supra* note 3.

Cooperation between AI companies can help mitigate these risks in a number of ways. Cooperation could speed up responsible AI development, for example, by pooling research on a socially beneficial application (e.g., image recognition for crop diseases in the Global South) or sharing details of a particularly good way to debias a model. Companies might agree to standards to help enhance AI safety, privacy, or guard against bias. Auditing mechanisms can help ensure that AI companies are keeping to their promises to develop systems that are safe, predictable, privacy-preserving, and non-biased. These kinds of cooperation benefit society and consumers. Cooperation strategies include agreements between competitors, mutual monitoring and information exchanges, and standard setting to proliferate best practices around an industry.³⁵ To promote trust building, companies should ideally implement these near-term strategies now, and sustain them over the long-term trajectory of TAI development.

Cooperation can also help in addressing longer-term harms if and when we approach TAI. Although companies will ideally agree to be bound by these agreements today, they may not be implemented until a future point in time. The Windfall Clause has been suggested as a way to secure an *ex ante* commitment from AI

³⁵ Note that in a forthcoming companion piece we will assess AI governance proposals that might intersect with abuse of dominance, public procurement and/or state aid concerns.

companies that they will redistribute their profits where those profits exceed a certain threshold. Second, the Assist Clause has been proposed to address the risk of AI companies “racing” to develop TAI first and underinvesting resources (such as researcher time or money) into system safety and reliability as a result.

We detail further the risks that each form of cooperation seeks to address in the sections analyzing the fourteen proposals.

II. BACKGROUND: THE PROHIBITION ON ANTI-COMPETITIVE AGREEMENTS UNDER ARTICLE 101

Competition law consists of rules that seek to protect the processes of market competition. The benefits of competition include lower prices, higher quality products, more innovation, less concentration of power, and greater efficiency than would be achieved under monopoly conditions. Consumer welfare, which focuses specifically on gains for consumers as opposed to society at large,³⁶ increases with more competition. Article 101(1) of the Treaty on the Functioning of the European Union (TFEU) is a key tool to preserve the rivalry and strategic uncertainty between competitors that drives the competitive dynamic, and prohibits any agreements that are capable of restricting competition. This section sets out a broad overview of Article 101(1) TFEU, the key

³⁶ The growth of Neo-Brandeisian or sustainability concerns in competition law is increasingly challenging the focus on consumers rather than society as a whole.

competition law provision that applies to proposed forms of cooperation between AI companies.³⁷

A. Article 101(1) Prohibition

Article 101(1) TFEU prohibits agreements that have an effect on trade between European Union countries and restrict competition in the EU. The text reads as follows:

“The following shall be prohibited as incompatible with the internal market: all agreements between undertakings, decisions by associations of undertakings and concerted practices which may affect trade between Member States and which have as their object or effect the prevention, restriction or distortion of competition within the internal market...”

The four key elements to a breach of Article 101(1) are: (1) there must be some form of agreement, decision, or concerted practice between undertakings; (2) which may affect trade between EU member states; (3) which has as its object or effect the restriction, prevention, or distortion of competition within the EU; and (4) which has an appreciable effect on competition. We explore the key elements of Article 101(1) below.

First, an analysis under Article 101(1) may require the definition of a “relevant market” on which effects on competition

³⁷ If the reader is already familiar with the basics of competition law, we invite them to skip to Section III (Legal Analysis of Proposed forms of Cooperation: Long-term Strategies).

are being assessed. This definition identifies the competitors of the undertakings involved who are capable of exerting competitive pressure on other undertakings.³⁸ The market is defined along two dimensions. The product market is the set of all products that a sufficient number of consumers would view as interchangeable.³⁹ The geographic market is the area of supply of the product where the competitive conditions are sufficiently homogenous, when compared to neighbouring areas.⁴⁰

A particular difficulty in the context of AI development is that the product market may be difficult to define with precision because product development is fast-moving and/or the relevant market does not yet exist (if the product is still at the R&D stage). Unlike in other areas of competition law (such as abuse of dominance), defining the relevant market on which the parties to an agreement operate is not a prerequisite to every Article 101(1) analysis.⁴¹ This paper will seek to define the relevant market only

³⁸ Commission Notice on the Definition of Relevant Market for the Purposes of Community Competition Law (97/C372/03), 1997 O.J. C 372 (“Market Definition Notice”), ¶ 2.

³⁹ This is so-called “demand-side substitutability” which is assessed by a hypothetical exercise that asks whether the parties’ customers would switch to readily available substitutes in response to a hypothetical small (in the range 5 % to 10 %) but permanent relative price increase in the products being considered. If substitution were enough to make the price increase unprofitable because of the resulting loss of sales, additional substitutes are included in the relevant market (*id.*, ¶ 17). Another factor is “supply-side substitutability,” which refers to the ease with which a supplier could switch its supply from one product to another.

⁴⁰ Market Definition Notice, *supra* note 38, at ¶¶ 7-9.

⁴¹ Case C-439/11, P Ziegler SA v Commission, EU:C:2013:513, at ¶ 63 (holding that it is unnecessary, in certain circumstances, to define the relevant market in order to determine if there is an appreciable effect on trade between Member

when it is necessary for the substantive analysis in each particular case.

Second, Article 101(1) only applies to an agreement between two or more independent “undertakings.” It does not apply, for example, to intra-group agreements i.e. agreements between entities within the same economic group.⁴² An undertaking can include any natural or legal person that is engaged in economic activity (that is, offering goods or services to a given market).⁴³ The activities of a non-profit organization can also be “economic activity” if the private sector provides them or could provide them—there need not be a profit-earning motive or economic purpose.⁴⁴ EU competition law is likely to treat an association of AI companies such as PAI as an undertaking under competition law.

Third, any form of “agreements, decisions, concerted practices” can fall within the scope of Article 101(1). Competition law is not concerned with the form of an anti-competitive agreement—otherwise, evasion of the law would be easy. The key

States for the purpose of Article 101(1), namely where it is possible to establish that the conduct in question is capable of affecting trade between Member States and has the object or effect of preventing, restricting or distorting competition, even in the absence of such a definition).

⁴² This includes for example an agreement between a parent company and its subsidiary. See e.g. Case 22/71- *Beguelin Import v. GL Import Export*, 1971 E.C.R. 949.

⁴³ Case C-41/90, *Hofner and Elser v Macrotron GmbH*, 1991 E.C.R. I-1979, ¶ 21; Joined cases C-180/98 to C-184/98, *Pavel Pavlov and Others v Stichting Pensioenfonds Medische Specialisten*, 2000 E.C.R. I-6451, ¶ 75.

⁴⁴ Joined cases 209 to 215 and 218/78, *Heintz van Landewyck SARL and others v Commission*, 1980 E.C.R. 3125, ¶ 88.

factor is that there is an expression of a joint intention or a “meeting of minds” to behave in a particular way between at least two undertakings.⁴⁵ Therefore, gentlemen's agreements, simple understandings and “protocols” have been held to be “agreements”.⁴⁶ “Concerted practices” refers to instances where businesses, although not entering into a binding agreement as such, “knowingly substitute practical cooperation between them for the risks of competition.”⁴⁷ The broad concept of an agreement covers the partial and conditional agreements in the negotiations that lead up to a final agreement.⁴⁸ If parties never implement an agreement, that is irrelevant to the existence of an agreement.⁴⁹ However, if an agreement is implemented, that could constitute a separate and additional breach of Article 101(1), or could be treated as part of a single and continuous infringement stemming from the original agreement.⁵⁰

⁴⁵ Case T-41/96, *Bayer v Commission*, 2000 E.C.R. III-3383, ¶ 69.

⁴⁶ Case 41/69, *ACF Chemiefarma NV v Commission*, 1970 E.C.R. 661 (*Quinine*); Case T-53/03, *BPB plc v Commission*, 2008 E.C.R. II-1333, ¶ 72 (*Plasterboard*); *Re Stichting Sigarettenindustrie Agreements* OJ [1982] L 232/1 (an ‘understanding’ between trade associations held to be an agreement). Case IV/29.525 and IV/30.000 - SSI, Commission decision of 15 July 1982, 1982 O.J. L 232.

⁴⁷ Case C-48/69, *Imperial Chemical Industries v European Commission (Dyestuffs)*, 1972 E.C.R. 619, ¶ 64.

⁴⁸ Case No IV/35.691/E-4, *Pre-Insulated Pipe Cartel*, Commission Decision of 21 October 1998, 1999 O.J. L24/1, ¶ 133, *upheld on appeal*, Case T-9/99, *HFB v Commission*, 2002 E.C.R. II-1487.

⁴⁹ Case COMP/39181, *Candle Waxes*, Commission decision of 1 October 2008, summary published in 2009 O.J. C295/17, ¶¶ 299-301, *upheld on appeal* Cases T-558/08, *Eni v Commission*, 2014 EU:T:2014:1080, ¶¶ 132-133.

⁵⁰ See *infra* Section III.A (Assist Clause as a “Single and Continuous Infringement”).

Article 101(1) contains a non-exhaustive list of examples of agreements that have as their “object or effect the prevention, restriction or distortion of competition” within the EU. The wording “object or effect” is significant. Agreements can have as their “object” the restriction of competition, wherein it is unnecessary to prove that it will provide anti-competitive effects. For an agreement to be a “by object” infringement, it must reveal, by its very nature, a sufficient degree of harm to competition that an analysis of its effects is not necessary to establish a breach of Article 101(1).⁵¹ Further, it tends to be difficult to argue that a “by object” infringement is exempt under Article 101(3), though it is theoretically possible to do so.⁵²

“By object” agreements between competitors include restrictions such as price fixing, output or sales limitations, paying competitors to delay the launch of competing products, and exchanging information that reduces uncertainty about competitor behavior. Only where it is unclear that an agreement has as its object the infringement of competition will it be necessary to consider whether it might have the effect of doing so. These are so-called “by

⁵¹ Case C-67/13 P, *Groupement des Cartes Bancaires v Commission*, EU:C:2014:1958, 12, ¶¶ 49 and 57.

⁵² See e.g. Case C-439/09, *Pierre Fabre Dermo-Cosmétique SAS v. Président de l’ Autorité de la concurrence and Ministre de l’ économie, de l’ industrie et de l’ emploi*, 2011 E.C.R. I- 9419, ¶¶ 49-57.

effect” infringements and require a more thorough assessment of the impact of the agreement on the market.

Anti-competitive agreements between competing (or potentially competing) parties at the same level of the supply chain are called “horizontal restrictions.” On the other hand, “vertical restrictions” are agreements between parties at different levels of the supply chain, for example a supplier and its reseller. Horizontal agreements are generally regarded as more serious breaches of Article 101(1) because they involve coordination between competitors and are therefore more likely to have a significant adverse impact on competition.⁵³

The analysis of this note will focus on horizontal agreements rather than vertical agreements, as those are the more common forms of cooperation currently discussed within AI governance.

B. Article 101(3) Exemption

Competition law is concerned with protecting competition to further consumer and societal welfare. If a company can demonstrate that the agreement in question can give rise to countervailing efficiencies that are passed on, then the effects of the

⁵³ The main reason why vertical restraints are generally seen as less harmful than horizontal restraints is because, in a horizontal relationship, a competitor that exercises its market power and increases its prices may benefit its competitors. This may incentivize competitors to induce each other to behave anti-competitively. In a vertical relationship, the product of one is the input for the other. Therefore companies in a vertical agreement usually wish to prevent the exercise of market power by the other (Guidelines on Vertical Restraints, 2010 O.J. (C 130) 1– 46 (“Vertical Guidelines”), ¶ 98).

Article 101(1) prohibition will be disapplied in relation to that agreement. In other words, Article 101(3) effectively acts as an exemption to Article 101(1).

To qualify for this legal exception, an agreement must satisfy each of the four conditions contained in Article 101(3). The burden of proof rests on the parties seeking to rely on the exemption.⁵⁴ The undertakings must put forward “convincing arguments and evidence” that the conditions of that provision are satisfied.⁵⁵

The conditions of the Article 101(3) exemption are as follows: (1) the agreement must contribute to improving the production or distribution of goods or to promoting technical or economic progress; (2) consumers must be allowed a fair share of the resulting benefit; (3) only restrictions indispensable to the achievement of those objectives can be imposed on the parties concerned; and (4) the parties should not be afforded the possibility of eliminating competition in respect of a substantial part of the products in question. We will look at each of these elements in turn.

⁵⁴ Article 2, Council Regulation (EC) No 1/2003 of 16 December 2002 on the implementation of the rules on competition laid down in Articles 81 and 82 of the Treaty (Text with EEA relevance), 2003 O.J. (L1) 1 (“Regulation 1/2003”). Note that under *Wouters*, the European Court of Justice found that a regulatory rule fell outside Article 101(1) altogether where ‘reasonably considered necessary in order to ensure the proper practice of the legal profession.’ (Case C-309/99, *Wouters v Algemene Raad van de Nederlandsche Orde van Advocaten*, 2002 E.C.R. I-1577) Although the circumstances that applied in *Wouters* were different to the Assist Clause (for example, in *Wouters* the offending agreement had a public law character), it highlights an alternative ‘route’ out of the Article 101(1) prohibition other than the Article 101(3) exemption. Similarly, see the ‘ancillary restraints’ analysis in discussion *infra* Section IV.B (Seconding staff (researchers or engineers)).

⁵⁵ Case C-68/12 *Slovenska sporitelna*, EU:C:2013:71, ¶ 32.

First, an agreement must produce efficiency gains that can be clearly substantiated and shown to flow directly from the agreement. This is a relatively high evidential standard and requires verification: the nature of the claimed efficiencies, the direct causal link between the agreement and the efficiencies; the likelihood and magnitude of each claimed efficiency, and how and when each claimed efficiency would be achieved. Generally, the parties are most likely to achieve efficiency gains if they combine complementary skills and assets, such as different research capabilities. In contrast, if the parties' skills and assets are substitutes rather than complements, an agreement is unlikely to lead to inefficiencies that can accrue to consumers.⁵⁶

Second, any restrictions on competition from the agreement should not go beyond what is necessary to achieve the efficiency gains. In other words, if there is a realistic alternative⁵⁷ that can achieve the same ethical or safety benefits without reducing the competition on the market, then the agreement will not satisfy Article 101(3). The parties must explain, for example, why they could not have achieved the same efficiencies by acting alone.⁵⁸

⁵⁶ Communication from the Commission — Notice — Guidelines on the application of Article 81(3) of the Treaty, 2004 O.J. (C 101) 97 (“Article 101(3) Guidelines”), ¶¶ 64-68.

⁵⁷ There must be no other economically practicable and less restrictive means of achieving the efficiencies, *id.*, ¶ 75.

⁵⁸ *Id.* at ¶ 76

Third, efficiency gains must be passed onto customers and outweigh the anti-competitive effects—the ethical or safety improvements should outweigh any negative consumer effects such as higher prices or lower output that result from the loss of competitive dynamics.⁵⁹

Whether efficiencies are passed on to consumers will depend on whether there is enough residual competition on the market and is closely linked to the fourth condition of Article 101(3), prohibiting the elimination of competition.⁶⁰ This condition requires an evaluation of the extent to which the agreement will reduce competition in the market. The EC considers both actual and potential competition when making this assessment.⁶¹ It examines the capacity of actual competitors to compete and their incentive to do so when assessing actual competition.⁶² The analysis of potential competition requires looking at barriers to entry facing undertakings that are not already competing within the relevant market, and whether there is the real possibility for new entry on a significant scale.⁶³

Article 101(3) is similar to the “Rule of Reason” under § 1 of the Sherman Act. Under § 1, agreements among competitors that

⁵⁹ *Id.* ¶ 85.

⁶⁰ *Id.* ¶¶ 95-97.

⁶¹ *Id.* ¶ 108.

⁶² *Id.* ¶¶ 109-10.

⁶³ *Id.* ¶¶ 114-15.

fix prices, allocate markets or customers, or restrict output are treated as per se unlawful. Other agreements are judged under the “Rule of Reason” when they have plausible efficiency justifications. The Rule of Reason weighs the potential anticompetitive effects of the agreement against its procompetitive benefits.

Competition law does not prohibit all forms of cooperation that restrict competition. For example, it recognizes certain categories of cooperation between competitors as beneficial for the parties and for consumer welfare. The EC acknowledges this in its Horizontal Cooperation Guidelines, noting that it “can be a means to share risk, save costs, increase investments, pool know-how, enhance product quality and variety, and launch innovation faster.” This paper will assess two of these types of cooperation, R&D cooperation and standardization agreements. Sections below will analyze how these forms of cooperation may be structured to reduce the risk of breaching Article 101(1): either because they do not have any restrictive effects (thereby falling outside Article 101(1) altogether), or because any efficiencies that they bring outweigh restrictive effect on competition, therefore engaging the Article 101(3) exemption. An agreement may benefit from the Article 101(3) exemption either through assessment of individual agreements on a case-by-case basis, or categories of agreements via one of the block exemptions. When an agreement fulfils the

conditions set out in a block exemption regulation, the agreement is automatically valid and enforceable.

C. Investigation and Enforcement Powers

Strategic cooperation between AI companies seeks to prevent a range of harms, from biased outputs to potentially catastrophic accidents, and is likely to involve large, well-resourced companies. Arguably, if the worst-case scenario is a fine from the EC, large companies might simply “price-in” this penalty as an unfortunate but necessary cost to achieve a far greater gain. It may follow that competition law cannot practically change the behaviour of these companies. However, this position is overly simplistic.

For one thing, the EC has the power to apply competition law to any company as long as the conduct has an effect on trade in the EU (called the “qualified effects” doctrine).⁶⁴ This extends to conduct or agreements entered into outside the EU if it has an effect within the EU, and covers companies that are not incorporated in the EU.⁶⁵ The string of competition law cases against US companies

⁶⁴ Case C-413/14 P, *Intel v Commission*, 2017 ECLI:EU:C:2017:632.

⁶⁵ *Id.* As well as the qualified effects doctrine, which requires there to be ‘immediate, substantial and foreseeable effects’ on within the EC, the EC can also establish jurisdiction on two other grounds- the single economic entity doctrine and the implementation doctrine, respectively. Under the single economic entity doctrine, a non-EU parent company of a group of companies can be held liable for an infringement committed by an EU subsidiary within the same group (Case C-48/69, *Imperial Chemical Industries v European Commission (Dyestuffs)*), *supra* note 47. Under the implementation doctrine, the decisive factor to establish jurisdiction is where the conduct was implemented not where it was entered into (Cases C-89/85, C-104/85, C-114/85, C-116/85, C-117/85 and C-125/85 to C-129/85, *A Ahlström Osakeyhtiö and others v European Commission*, 1988 EU:C:1988:447 (*Wood Pulp*)).

like Google, Qualcomm, and Microsoft, and ongoing cases against Apple and Amazon,⁶⁶ demonstrate this.⁶⁷ Another way for the EC to establish subject matter jurisdiction is by using an AI company's EU subsidiary as an "anchor."⁶⁸

Even if AI cooperation raises competition concerns in principle, in practice the EC might never find out about it. This is unlikely, however. The EC uses market intelligence monitoring to identify potential issues and it is likely to focus particularly on digital markets and the AI development space in particular for the foreseeable future.⁶⁹

Another important avenue for identifying breaches is through complaints from either a customer or competitor. A disgruntled competitor who has fallen behind in the lucrative AI development race could bring a claim against those leading the race in order to clip their wings. That disgruntled competitor may have

⁶⁶ See, e.g., Case AT.39740, Google Search (Shopping), *supra* note 10; Case AT.40099, Google Android, *supra* note 10; Case AT.40411, Google Search (AdSense), *supra* note 10; Case COMP/C-3/37.792, Microsoft, *supra* note 10; Case AT.39711, Qualcomm (predation), Commission decision of 18 July 2019; Case AT.40220, Qualcomm (exclusivity payments), Commission decision of 24 January 2018, 2018 O.J. C269/25.

⁶⁷ Under the economic entity doctrine, the conduct of a subsidiary active in the EC is attributed for antitrust purposes to the parent company seated outside the EC but exercising its corporate control on the subsidiary (Case C-48/69, Imperial Chemical Industries v European Commission (Dyestuffs), *supra* note 71, at ¶¶ 125-141.

⁶⁸ *Wood Pulp*, A Ahlström Osakeyhtiö and others v European Commission, *supra* note 65, at ¶¶ 11-23. The question of extraterritorial application of EU competition law, including the power of the EC to compel compliance or punish non-compliance with its laws (enforcement jurisdiction), is further discussed in our forthcoming paper on EU competition law.

⁶⁹ See discussion *infra* Section II.B (Investigation and Enforcement Powers).

been party to some of the discussions leading up to forms of cooperation, and therefore be able to provide insightful evidence to a regulator.⁷⁰ Furthermore, the pressure that a well-resourced complainant can place on regulators to bring an investigation against their rivals should not be underestimated. In many cases, powerful complainants are the driving force behind EC investigations, including the *Google* and *Microsoft* cases.

As mentioned above, the regulation of digital markets and Big Tech and the development of AI in Europe are currently hot topics in the EU. This focus suggests that any conduct between AI companies, especially if they happen to be one of the Big Tech companies, may be heavily scrutinized. If it identifies potential issues, the EC is likely to prioritise the investigation of these cases over others. Furthermore, the uncertainty hanging over companies during the course of investigations can have a detrimental effect on their incentives and ability to innovate. The length of investigations exacerbates this uncertainty—for example, the *Google Search* investigation lasted for 8 years.⁷¹ Competition law investigations

⁷⁰ Complainants are also entitled to certain rights to participate in any investigation that may arise from their complaint, including providing and receiving evidence on the file, as well as to potentially participate in any oral hearing. (Commission Regulation (EC) No 773/2004 of 7 April 2004 relating to the conduct of proceedings by the Commission pursuant to Articles 81 and 82 of the EC Treaty, 2004 O.J. (L 123), Article 6(1) & 6(2)).

⁷¹ Case AT.40411, *Google AdSense*, *supra* note 10; see also Damian Reece and Stephen Castle, *Microsoft rivals line up to sue after EU ruling*, THE INDEPENDENT (March 25, 2004), <https://www.independent.co.uk/news/business/news/microsoft-rivals-line-up-to-sue-after-eu-ruling-756881.html>.

against companies that are active in digital markets, such as AI, may be more likely to see longer competition law investigations. This is because the complexity and fast-moving nature of the industry makes it more difficult for regulators to make their competition assessment. As a result, regulators may seek more in-depth information from the business in order to make their assessment.

Serious consequences flow from competition law enforcement. For an intentional or negligent breach of Article 101, the EC can impose a fine of up to 10% of annual global turnover. The agreement will also be void and unenforceable. Examples of fines include EUR 4.3bn for Google's *Android* case in 2018, EUR 997m for Qualcomm in 2018 for its case on exclusivity payments and EUR 1.06bn for Intel in 2009. These are in addition to the time and resources necessary to defend long, in-depth competition cases.

However, this does not mean that a well-resourced AI company could simply choose to pay the fine and carry on with its unlawful cooperation strategy. Aside from the fine, the EC has the power to order the termination of the offending conduct. It can make an order for cessation of the infringing conduct, as well as an order requiring *inter alia* the making of supplies to other parties on particular terms,⁷² the suspension of a proposed merger pending

⁷² *C.f.* the *Google Search* case, where the EC avoided specifying precisely how Google should remedy its abuse, confining itself instead to ordering that the conduct be terminated and requiring Google to submit evidence of compliance

investigation, and the divestment of assets⁷³ (including the separation of previously combined assets).⁷⁴

If the EC makes an order and the company fails to follow it, the EC can impose periodic penalty payments of up to 5% of the average daily turnover in the preceding business year in order to enforce its order.⁷⁵ Although the EC has used this power relatively rarely, in 2006 and 2008 it imposed a daily penalty of EUR 1.5 million on Microsoft (totalling EUR 280.5 million) for failure to comply with an obligation to provide interoperability information contained in a decision from 2004.⁷⁶

The EC can impose interim measures on companies prior to the finding of infringement. This allows the EC to order the company under investigation to stop certain behavior, during the course of the investigation, if it deems that there is the risk of serious and irreparable damage to competition caused by that company.⁷⁷

Third parties harmed by breaches of EU competition law can also

within 90 days of the decision or face daily penalties of up to 5% of global group turnover. See Case AT.39740, *Google Shopping*, *supra* note 10, at Recital 700, summarizing the obligations and Pinar Akman, *The Theory of Abuse in Google Search: A Positive and Normative Assessment Under EU Competition Law*, 2017 J. L. TECH. & POLICY 365.

⁷³ However, the EC may only impose a divestment remedy in limited circumstances (i.e., where other, less interventionist remedies are not sufficient.).

⁷⁴ Article 7, Regulation 1/2003, *supra* note 54.

⁷⁵ *Id.*, Article 24.

⁷⁶ This amount was later reduced by EUR 39million on appeal to the General Court (Case T-167/08, *Microsoft Corp v European Commission*, judgment of 27 June 2012).

⁷⁷ Regulation 1/2003, *supra* note 54, Article 8. Though rarely used, interim measures were imposed on Broadcom in October 2019 and the EC has expressed a desire to use them more frequently in digital markets in future.

bring a private damages action in their domestic courts. These can include private action claims seeking large sums in damages sums.⁷⁸

This paper analyzes the state of competition law as at the date of writing. However, EU competition law is going through a series of important changes. These include the impact of the COVID-19 crisis, Brexit, and the focus on regulating digital markets and Big Tech. A detailed assessment of these trends and events are outside the scope of this paper, but we draw out a few of the key themes and implications below.⁷⁹

The EC has highlighted that even during the exceptional circumstances of COVID-19, it will continue to closely monitor the market for any breaches of EU competition law. However, the EC has published a communication⁸⁰ that permits certain forms of cooperation between companies and trade associations that would usually fall foul of competition rules⁸¹ if they are necessary to ensure

⁷⁸ For example, Barclays, Royal Bank of Scotland and three other banks are currently being sued by investors for at least £1 billion over rigging of the foreign exchange market in a class action in the UK (Sean Farrell, *Barclays, RBS and Other Banks Face £1bn Forex Rigging Lawsuit*, THE GUARDIAN (Jul. 29, 2019), <https://www.theguardian.com/business/2019/jul/29/barclays-rbs-banks-forex-rigging-lawsuit-jp-morgan-citigroup-ubs>). In another case, retailers including Asda and Argos “are in line for potential billion-pound payouts after the UK’s highest court ruled that transaction fees charged by Visa and Mastercard breached competition laws” (Jane Croft, *Mastercard and Visa face billion-pound payouts after UK court ruling*, FT (Jun. 17, 2020), <https://www.ft.com/content/e948aa14-0f44-4e47-ad2e-397452d859c5>).

⁷⁹ Some of these will be explored in our forthcoming paper, in particular the regulation of big tech, see *supra* note 35.

⁸⁰ Temporary Framework for assessing antitrust issues related to business cooperation in response to situations of urgency stemming from the current COVID-19 outbreak, Apr. 8, 2020, 2020 O.J. (C 116 I) 02.

⁸¹ This includes measures to adapt production, stock management and possibly exchange of information related to the distribution and production of certain medicines.

the supply and adequate distribution of essential and scarce products and comply with certain conditions.

The EC's communication clearly applies to the exceptional circumstances of COVID-19. However, this might introduce a broader or more flexible approach, for example to the application of the exemption under Article 101(3) for countervailing efficiencies, allowing more forms of AI cooperation to comply with competition law.

The UK left the EU on January 31, 2020. After that period, there are unlikely to be major changes that will affect the assessment of forms of AI cooperation we analyze in this paper under the Article 101(1) regime.⁸² Any cooperation and agreements between UK-based or incorporated AI companies will still be within the jurisdiction of the EU competition rules and must comply with Article 101(1) if they have an effect within the EU ("effect on trade" rule).⁸³

Where an agreement does not affect trade within the EU, for example if it only has localized effects on the UK market, the UK competition authority will apply the UK equivalent of Article 101(1), called the Chapter I prohibition. The substance of the

⁸² Brexit will have more significant impact on state aid rules, however that will be assessed under our separate forthcoming paper, see *supra* note 35.

⁸³ This is in the same way that Asian and US businesses are subject to EU competition law where their agreements and conduct affect EU markets. For example, a UK participant in a global cartel will continue to face investigation and fines by the EC.

Chapter I prohibition is very similar to Article 101(1), but over time it is possible the UK courts will diverge in their interpretation of the prohibition from the EU jurisprudence. This is because they will no longer be subject to a duty to interpret the UK rules in a manner consistent with competition case law of the European Court of Justice.⁸⁴

In addition, one of the current hot topics in EU competition law—and competition law globally⁸⁵—is the regulation of digital markets and Big Tech. The EC in its Experts report, for example, considered whether competition law should be strengthened or

⁸⁴ Section 60 of the Competition Act 1998 (that contains the duty to interpret the UK rules consistently with EU caselaw) is revoked by the Competition SI (Competition (Amendment etc.) (EU Exit) Regulations 2019, SI 1993 No. 93) and in its place a new provision, section 60A, is inserted with effect from Brexit. The new section 60A merely requires the CMA and courts to avoid inconsistency between their decisions and EU law and the decisions of the European Court of Justice before exit day. In addition, in the draft Brexit trade deal (UK-EU Trade and Cooperation Agreement, dated 28 December 2020 accessible here: https://ec.europa.eu/info/sites/info/files/brexit_files/info_site/tca-20-12-28.pdf), both the EU and UK commit to “maintain their high standards of competition law” to “effectively [address] anticompetitive practices [including] agreements between economic actors, decisions by associations of economic actors and concerted practices which have as their object or effect the prevention, restriction or distortion of competition.” This mirrors the wording of Article 101(1). *See* UK-EU Trade and Cooperation Agreement, Part 2 (Trade, Transport, Fisheries and Other Arrangements), Title XI (Level Playing Field for Open and Fair Competition and Sustainable Development), Chapter 2 (Competition Policy), Article 2.2 (Competition Law), Sub-article 1(a).

⁸⁵ See also similar proposals by the regulators in the UK, France, Germany, USA, Australia, and others, as well as recent joint proposed options for modernising EU competition law by France, Germany and Poland, which include increased scrutiny of Big Tech, as well as proposals to reform merger control rules to facilitate the formation of European national champions. The overall effect on forms of AI cooperation is heightened scrutiny of AI companies, but possibly more relaxed merger control rules if the AI companies merging are European. For further discussion, *see* Joint Statement by France, Germany and Poland, *Modernising EU Competition Law* (Jul. 4, 2019), https://www.bmwi.de/Redaktion/DE/Downloads/M-O/modernising-eu-competition-policy.pdf?__blob=publicationFile&v=4.

reformed to sufficiently scrutinize digital markets and platforms.⁸⁶

We offer a fuller analysis of the implications of competition law's rise against Big Tech in a separate paper. However, the overall effect is likely to be an increased scrutiny of the AI sector.⁸⁷

This goes hand in hand with the EU's central strategy of encouraging AI development in Europe, as set out in the AI White Paper. In that paper, the EC highlights the importance of data and in training machine learning systems, the "accumulation of vast amounts of data" by Big Tech companies, and the market imbalances in relation to the access and use of data by SMEs. This strategic focus on AI and big tech's position of strength in this field will probably only increase with advances in AI development, and may result in a heightened level of scrutiny over all AI companies and over Big Tech's activities in AI development.⁸⁸

For example, the EC is currently considering an ex ante regulatory regime for markets characterized by large digital platforms that act as gatekeepers to ensure the markets "stay fair and open." On December 15, 2020, the EC published a proposed regulation called the Digital Markets Act (DMA) on contestable and

⁸⁶ See e.g. JACQUES CRÉMER, YVES-ALEXANDRE DE MONTJOYE AND HEIKE SCHWEITZER, SHAPING COMPETITION POLICY FOR THE DIGITAL ERA, EC REPORT (4 April 2019), <http://ec.europa.eu/competition/publications/reports/kd0419345enn.pdf>.

⁸⁷ *Id.*

⁸⁸ See European Commission Communication 'Shaping Europe's Digital Future', COM 67 (Feb. 19, 2020), https://ec.europa.eu/info/sites/info/files/communication-shaping-europes-digital-future-feb2020_en_4.pdf.

fair markets in the digital sector.⁸⁹ This regulation targets the negative consequences arising from certain behaviors by platforms acting as designated digital “gatekeepers” to the single market. The proposed DMA addresses unfair practices by gatekeepers that either fall outside the existing EU competition control rules, or which these rules cannot always effectively tackle because of the systemic nature of some behaviors, as well as the ex-post and case-by-case nature of competition law. The powers in the DMA are additional and complementary to existing competition enforcement tools.

The proposed DMA contains a series of obligations and prohibitions relating to self-preferencing, interoperability, data-related practices and tying. The DMA Proposal gives the EC new enforcement and sanctioning powers, including a new power to carry out targeted market investigations. Specifically, the Commission would be empowered to conduct investigations to identify designated gatekeepers and their core services, and to ensure their compliance with the DMA.

It is still too early to know what the regime will eventually look like and understand its full implications, but it will undoubtedly add an additional layer of scrutiny to the conduct of AI companies.

⁸⁹ Proposal For A Regulation Of The European Parliament And Of The Council On Contestable And Fair Markets In The Digital Sector (Digital Markets Act), COM/2020/0842 (Dec. 15, 2020).

III. LEGAL ANALYSIS OF PROPOSED FORMS OF COOPERATION: LONG-TERM STRATEGIES

In the next two sections, we survey two broad types of cooperation that are under consideration within AI governance. In this Section, we consider the Assist Clause and the Windfall Clause. Both strategies are designed to address risks in a scenario where TAI is developed or nearing development.

Even though implementation of these strategies may not occur until some point in the future, competition law concerns can arise as soon as the companies agree to be bound, regardless of if or when the strategies are ever implemented. For both the Assist Clause and Windfall Clause, we explain the relevant strategies, what useful purposes they would serve, what competition concerns they might raise, and how to address those concerns. In the following Section, we analyze other strategies to encourage safe and responsible AI development, including mutual or third-party monitoring, standard-setting, and strategic mergers

A preliminary question is whether the EC will have jurisdiction to apply EU competition law to, for example, a US AI company that engages in an anti-competitive agreement with another US company. We do not propose to look at whether jurisdiction is established in every case, but we generally posit that the EC's jurisdiction is likely to be established. This is because (1) a market for AI products or services is likely to be global, so any

agreement between AI companies is likely to have direct or knock-on effects on EU trade; and (2) the AI companies that are the focus of our discussion are likely to be large, global companies with European subsidiaries.⁹⁰

A. OpenAI Assist Clause

Key Recommendations

There is a risk that AI companies binding themselves to an Assist Clause similar to OpenAI's could breach Article 101(1). The Assist Clause is essentially a commitment from one AI company not to compete with another. If the Assist Clause leads to AI companies who are close competitors agreeing not to compete with each other, this could lead to a significant reduction in competition given the otherwise fierce competition between those companies that is now lost. A company could breach competition law as soon as it reached an agreement with another company to be bound by the Assist Clause, regardless of when or if it is implemented or triggered.

To mitigate this risk, OpenAI (and any other AI company) should take care not to contact competitors to persuade them to follow the Assist Clause, or otherwise seek to influence or reach an understanding with competitors to that effect. This will help ensure that any company's decision to be bound by an Assist Clause is truly unilateral, thus falling outside of the Article 101(1) prohibition.

⁹⁰ As discussed in *infra* Section II.B (Investigation and Enforcement Powers), the EC's jurisdiction can be founded on both these grounds.

Separately, implementation of the Assist Clause through a competitor assisting the Leader can also constitute a separate breach of Article 101(1), or form part of a single and continuous infringement stemming from the original plan, that is the Assist Clause itself.

To mitigate risks from the Assist Clause's implementation, we advise that cooperation occur at as early a stage as possible. Cooperation should also be widespread enough across the TAI development industry to dampen harmful race dynamics, but still leave sufficient competition in the market to mitigate competition law risk.

Amongst the different ways to implement the Assist Clause, a company can significantly mitigate risks if cooperation between a competitor and the leading TAI developer does not lead to the loss or shutdown of the competitor's AI development activities. It can achieve this by *inter alia* shifting R&D and production to AI safety rather than on product development.

Analysis

A central risk from TAI development is that it may result in harmful race dynamics. Given the great rewards that may accrue to the first to achieve TAI, the fear is that this would lead to a frantic development race that leads to underinvestment of resources (researcher time or money) into system safety and reliability.

In order to address these harmful race dynamics, OpenAI proposed in April 2018 the “Assist Clause” in its Charter. This clause states that:

We are concerned about late-stage AGI development becoming a competitive race without time for adequate safety precautions. Therefore, if a value-aligned, safety-conscious project comes close to building AGI before we do, we commit to stop competing with and start assisting this project. We will work out specifics in case-by-case agreements, but a typical triggering condition might be “a better-than-even chance of success in the next two years.”⁹¹

The following sections analyze the compatibility of the Assist Clause with Article 101(1) in two steps. First, we consider whether the Assist Clause itself could be a restrictive agreement in breach of Article 101(1). Next, we look at whether different forms of implementing the Assist Clause could constitute an anti-competitive agreement or form of cooperation that also breaches Article 101(1).

This two-step analysis is necessary because the Assist Clause could breach Article 101(1) even if partially or never implemented.⁹² A “concurrence of wills” alone is enough to trigger a breach of Article 101(1). If the Assist Clause does not breach Article 101(1), its subsequent implementation (for example, through

⁹¹ OpenAI Charter, *supra* note 22.

⁹² Case T-558/08, *Eni v Commission*, *supra* note 49, ¶¶ 132-133 (General Court held that it was sufficient to show a concurrence of wills on price-fixing; it was not necessary to show that price increases were, or could be, implemented).

an R&D cooperation agreement) could still constitute a separate breach. Alternatively, both the Assist Clause and its subsequent implementation may breach Article 101(1). The EC could treat these two forms of conduct as a “single and continuous infringement”, which has implications for the duration of the infringement and therefore the level of fine, as well as for the parties that could be held liable for the infringement.⁹³

Preliminary Question: Could the Assist Clause itself be an agreement between undertakings?

OpenAI’s Assist Clause is currently a public, unilateral statement in its Charter, published on their website. If OpenAI’s decision to be bound by the Clause is truly unilateral in nature, it will fall outside of Article 101(1), which applies only to agreements between undertakings.⁹⁴ However, as mentioned above, the concept of an anti-competitive agreement under Article 101(1) can be very wide. For example, if OpenAI’s competitors publicly announce their own Assist Clauses, EU competition law could deem the responses of competitors to each other’s public announcements a strategy for reaching a common understanding on future commercial behavior.⁹⁵

⁹³ See *infra* Section III.A (Assist Clause as a ‘single and continuous infringement’).

⁹⁴ Communication from the Commission - Guidelines on the applicability of Article 101 of the Treaty on the Functioning of the European Union to horizontal co-operation agreements, 2011 O.J. (C11) 1 (“Horizontal Cooperation Guidelines”), ¶ 63.

⁹⁵ *Id.*

The EC has held that this type of ostensibly unilateral, public announcement can constitute an agreement under Article 101(1) where it reduces the level of uncertainty about a competitor's future commercial behavior.⁹⁶ Risks could arise, therefore, if regulators see the Assist Clause as an invitation to other AI companies to coordinate their behavior. OpenAI (and any other companies that adhere to the Assist Clause) (“Assist Clause participants”) should take care not to contact competitors to persuade them to follow the Assist Clause, or otherwise seek to influence or reach an understanding with competitors to that effect.⁹⁷

Could an agreement that implements the Assist Clause breach Article 101(1)?

In the previous section, we looked at whether the Assist Clause itself could be an agreement that falls within Article 101(1). However, even if competition law does not find the Assist Clause itself to be an agreement under Article 101(1), this section will

⁹⁶ In the 2016 *Liner Shipping* investigation, for example, the Commission reached the preliminary view that announcing future price increases breached Article 101(1). This is despite the fact that there was no direct agreement or even covert contact between the parties. The public announcements may signal the intended market conduct of carriers and decrease their incentives to compete against each other (Case COMP/39850, *Container Shipping*, Commission Decision of 31 August 2016).

⁹⁷ At the same time, Article 101 does not “deprive companies of the right to adapt themselves intelligently to the existing or anticipated conduct of their competitors” (Horizontal Cooperation Guidelines, *supra* note 94, ¶ 61). Therefore, it does not prohibit an AI company adapting its behavior as a reaction to / influenced by the OpenAI clause, as long as it is a genuinely unilateral decision.

consider the question of whether the implementation of the Assist Clause could constitute a separate breach of Article 101(1).

For present purposes, we will not consider in detail the alternative scenario where both the Assist Clause itself and its implementation breach Article 101(1) and form a “single and continuous infringement.”

There are several ways that the Assist Clause could be implemented by a competitor to assist the leading TAI developer (“Leader”). Each of them could also constitute, in addition to the Assist Clause, a restrictive agreement between competitors in breach of Article 101(1).

The methods of implementation largely fall into two categories: first, the competitor assists the Leader with the latter’s TAI development, including through a JV or by merging with the Leader (“collaborative implementation”); or second, the competitor slows down TAI development (for instance, by firing employees) or switches resources from TAI development to another area such as safety research or applied research. In the second category, there is no direct cooperation or merger with the Leader; instead, implementation is unilateral (“non-collaborative implementation”).

Competition law could classify collaborative implementation scenarios as a horizontal cooperation agreement on R&D. In addition, depending on the way the parties cooperate, EU

merger approval may also be required, though this falls outside the scope of this paper. Non-collaborative implementation scenarios will fall outside Article 101(1) entirely if they are truly unilateral based on the same principles as discussed in the previous section.⁹⁸ We assess both categories of implementation under Article 101(1) below.

*Category 1: Collaborative implementation of Assist Clause:
R&D Cooperation*

The following analysis will first consider whether collaborative implementation of the Assist clause through R&D cooperation falls within the Article 101(1) prohibition, and if so, whether it can benefit from the R&D block exemption that sets out conditions under which forms of R&D cooperation can be exempt from Article 101(1). Failing that, we will consider whether the R&D cooperation is likely to breach Article 101(1) under the framework of the Horizontal Guidelines, and if so whether the cooperation can benefit from the individual exemption under Article 101(3).

Defining the market is an important first step in the analysis of R&D agreements because it allows us to assess whether collaborative implementation of the Assist clause through R&D cooperation falls within the Article 101(1) prohibition.⁹⁹

⁹⁸ The relevant questions are whether with the existence of the Assist Clause in the OpenAI Charter, and possibly others in the industry at the relevant point in time, it can truly be said that the competitor's acts are unilateral.

⁹⁹ The Commission Regulation (EU) No 1217/2010 of 14 December 2010 on the application of Article 101(3) of the Treaty on the Functioning of the European

The market definition will depend on whether we see TAI as creating an entirely new technology or as improving existing products or technologies, which may in turn depend on how gradual the development trajectory is. It seems likely that the R&D cooperation will lie somewhere between the two scenarios of developing an entirely new technology or improving existing technologies. We should thus assess TAI development on both frameworks of analysis described below: as an innovation market and as an existing technology market.¹⁰⁰

A more discontinuous development trajectory will more likely lead to “the development of new products or technology which either may—if emerging—one day replace existing ones or which are being developed for a new intended use and will therefore not replace existing products but create a completely new demand.”¹⁰¹ Under this scenario, we should assess the R&D coordination according to its effects on competition in innovation for TAI, rather than on existing market shares.¹⁰²

It is also possible that the development of TAI is incremental, such that we can see it as a significant improvement to

Union to certain categories of research and development agreements, 2010 O.J. (L 335) 36 (“R&D Block Exemption”) defines categories of R&D agreements that the Commission views as normally satisfying the Article 101(3) exemption (*see* R&D Block Exemption, ¶ 2).

¹⁰⁰ This “blended” approach is discussed in the Horizontal Cooperation Guidelines, *supra* note 94, ¶¶ 112 & 139.

¹⁰¹ *Id.* at ¶ 119.

¹⁰² As discussed further in the section below ‘Analysis of R&D Cooperation under Article 101(1)’.

an existing technology, rather than a technology for an entirely new intended use and creating completely new demand. In this scenario, in the competition analysis we should account for the competitor's and leader's market shares on the existing technology market that TAI is improving or eventually replacing.¹⁰³

There is still considerable uncertainty as to which of the above two scenarios TAI development will fall under.¹⁰⁴ The following analysis is also relevant to the possibility that it may lie somewhere between the two; that is, where innovation efforts may create technology that, over time, replaces existing ones. The following analysis will therefore assess competitive conditions on both relevant markets.¹⁰⁵

Having defined the relevant market, the next step is to assess whether the R&D cooperation could restrict competition under Article 101(1) or if it falls outside Article 101(1) altogether. Without carrying out a full assessment, we can preliminarily assess that there is a risk that R&D cooperation is prohibited under Article 101(1). The Leader is likely to have a strong market position in existing and related markets to TAI. In addition, the company intends the R&D cooperation to reduce or slow down innovation.¹⁰⁶ Having

¹⁰³ *Id.* at ¶¶ 116-118.

¹⁰⁴ *See, e.g.*, Gruetzemacher & Whittlestone, *supra* note 21.

¹⁰⁵ Horizontal Cooperation Guidelines, *supra* note 94, ¶ 112.

¹⁰⁶ The Horizontal Cooperation Guidelines explain that R&D cooperation agreements could lead to possible anti-competitive effects, such as the reduction or slowing down of innovation, such that fewer or worse products come to the market or do so later than they otherwise would.

established that there is a risk of breaching Article 101(1) that warrants continuing the assessment under Article 101(1), we will next consider whether the R&D cooperation can benefit from the “safe harbor” under the R&D Block Exemption.¹⁰⁷

Is the R&D cooperation exempt under the R&D Block Exemption?

To benefit from the R&D Block Exemption, the R&D cooperation must: (1) be a qualifying R&D cooperation agreement as defined in the R&D Block Exemption; (2) satisfy the market share thresholds on the existing technology market and conditions on the innovation market; and (3) not contain any blacklisted “hardcore restrictions.”

The R&D Block Exemption definition of an R&D cooperation agreement includes any agreement between parties relating to the conditions under which they pursue joint R&D of contract¹⁰⁸ products and technologies.¹⁰⁹ For R&D to be “joint,” it can be: (1) carried out by a joint team or undertaking; (2) jointly entrusted to a third party; or (3) allocated between the parties in any way they consider most appropriate.¹¹⁰ If the company implements

¹⁰⁷ R&D Block Exemption, *supra* note 99, Article 2(1). Note that the current version of the R&D Block Exemption is in place until December 31st, 2022. The EC is currently consulting on the effectiveness, efficiency, relevance, coherence and EU added value of the R&D Block Exemption to determine whether to allow it to lapse, prolong its duration, or revise it and the Horizontal Cooperation Guidelines to take proper account of new market developments since 2010.

¹⁰⁸ “Contract” products and technologies refer to those arising out of the joint R&D or manufactured applying the contract technologies. *Id.* at Article 1(1)(e).

¹⁰⁹ *Id.* at Article 1(1)(a).

¹¹⁰ *Id.* at Article 1(1)(m) & (1)(n).

the Assist Clause in any of these “joint” ways, it will fall within the scope of the R&D Block Exemption.¹¹¹

Where the R&D cooperation develops a product which creates a completely new demand, the R&D Block Exemption treats those agreements as benefiting from the Article 101(3) exemption for the duration of R&D and for seven further years after the product first reaches the market.¹¹² The agreement can benefit from the exemption if the following conditions are satisfied. The R&D cooperation must not:¹¹³ (1) restrict the freedom of the parties to carry out R&D in an unrelated field, such as outside TAI development¹¹⁴; (2) limit output or sales¹¹⁵; or (3) “eliminate[] effective competition in innovation.”¹¹⁶ If any of these apply, the R&D cooperation will not benefit from the R&D Block Exemption.

Where TAI development is more incremental, existing technology markets could be significantly improved or replaced by

¹¹¹ Note that the definition of R&D agreement under the R&D Block Exemption and the Horizontal Cooperation Guidelines are different. The definition under the R&D is narrower.

¹¹² Horizontal Cooperation Guidelines, *supra* note 94, ¶ 126; R&D Block Exemption, *supra* note 99, Article 4(1).

¹¹³ For full list of conditions, see Article 3 (conditions for exemption), Article 5 (hardcore restrictions (any R&D agreement that contains a hardcore restriction will not be able to rely on the R&D Block Exemption at all) and Article 6 (excluded restrictions, i.e. restrictions within an R&D agreement that cannot be exempt under the R&D Block Exemption, but the rest of the agreement may still be exempt under the R&D Block Exemption).

¹¹⁴ R&D Block Exemption, *supra* note 99, Article 5(a).

¹¹⁵ *Id.* at Article 5(b).

¹¹⁶ This “elimination of competition” scenario may occur if R&D cooperation applies across large swathes of the TAI development market i.e. all companies that are credibly developing, or could credibly develop, TAI are cooperating with the Leader rather than competing with each other. *See Id.* at Articles 19-21; Horizontal Cooperation Guidelines, *supra* note 94, ¶ 126.

TAI rather than the latter creating a completely new market altogether. In this scenario, the parties' positions on the existing technology market are relevant to the analysis. This is because we can deem that competitors on the existing technology market competitively constrain the TAI market to some degree.

To benefit from the R&D Block Exemption, the parties' combined market shares of licensing fees from those existing technologies, and the share of downstream sales of products or services incorporating the licensed technologies, cannot exceed 25% on either of these markets.¹¹⁷ Given the likelihood that the Leader will also have a strong market position in TAI's predecessor technologies alone, the parties might exceed this market share cap and therefore not benefit from the R&D Block Exemption.

In conclusion, the parties will benefit from the R&D Block Exemption if: (1) their combined market share is below the 25% market share threshold; (2) they satisfy the conditions outlined above with respect to assessment under the innovation market (including no limitation of output or sales and no elimination of competition); and (3) the R&D agreement does not contain any other "hardcore" restrictions on competition.¹¹⁸ Outside of the R&D

¹¹⁷ Horizontal Cooperation Guidelines, *supra* note 94, ¶ 125.

¹¹⁸ E.g., restriction on the parties' ability to carry out R&D independently or with third parties in unconnected fields, the limitation of output or sales (subject to some exceptions), the fixing of prices when selling the product to third parties (subject to some exceptions). For the full list of hardcore restrictions, see Article 5 of the R&D Block Exemption, *supra* note 83. Note that if an agreement contains

Block Exemption, competition law would assess the R&D cooperation under the usual Article 101(1) principles, to which we now turn.¹¹⁹

Analysis of R&D Cooperation under Article 101(1)

The type of cooperation envisioned under the Assist Clause is likely to qualify as an R&D agreement within the scope of the Horizontal Cooperation Guidelines because: (1) it is an agreement between competitors (2) to jointly improve existing technologies or carry out R&D on completely new products and (3) the form of collaboration could be in any form, such as via a cooperation agreement or a jointly controlled company.¹²⁰

The Horizontal Cooperation Guidelines explain that possible anti-competitive effects of R&D cooperation include the reduction or slowing down of innovation, such that fewer or worse products come to the market or do so later than they otherwise would.¹²¹ These restrictive effects on competition are only likely where the competitor and/or the Leader have market power on the existing markets, and/or the R&D cooperation appreciably reduces competition in innovation.¹²²

hardcore restrictions, it will be entirely excluded from the benefit of the block exemption.

¹¹⁹ Horizontal Cooperation Guidelines, *supra* note 94, ¶ 135.

¹²⁰ *Id.* ¶ 111.

¹²¹ *Id.* ¶ 127.

¹²² *Id.* ¶ 127.

If TAI development produces an entirely new technology that creates its own market, the legal analysis would focus on the potential for R&D cooperation to restrict innovation.¹²³ Under this analysis, R&D cooperation would likely be problematic where: (1) the competitor and Leader could both have independently developed TAI on their own; (2) they cooperated at a relatively late stage of development, such that they were “rather near” the launch of the product, and (3) the remaining firms racing to develop TAI on the market are not viable competitors.¹²⁴

On the other hand, competition concerns are less likely to arise if the coordinating parties did not have the means to carry out the necessary R&D independently but are able to do so jointly because they are bringing together these complementary resources. There is no loss of competition as a result of the cooperation because the coordinating parties would not have been able to effectively compete with each other whilst acting independently.¹²⁵

It follows that, to mitigate competition law risks, it would be advisable that: (1) cooperation occurs at as early a stage as possible and (2) cooperation with the Leader is widespread enough across the TAI development industry to dampen harmful race dynamics, but

¹²³ *Id.* ¶ 138.

¹²⁴ *Id.* ¶¶ 119-122.

¹²⁵ *Id.* at ¶ 130.

still leaves sufficient competition on the market to satisfy competition law.

R&D cooperation would be problematic under Article 101(1) if the parties have a strong position in the existing technology market and face little competitive constraint there. This would be even more of a concern if the parties also hold a strong position in the innovation market. Several factors are relevant in assessing whether R&D cooperation restricts competition under Article 101(1) on an existing technology market. First, the market positions of the parties are relevant. The R&D cooperation may breach Article 101(1) if it involves significant competitors on an existing technology market who cooperate to develop a new technology that may one day replace the existing (but related) technology. A combined market share of above roughly 40-50% may start to indicate competition concerns¹²⁶ although this will vary depending on various other factors.

Second, there should be an assessment of the competition on the market. The R&D cooperation may be problematic if the competitors are each other's closest competitor in the market, and the remaining competitors are not sufficiently viable to constitute a competitive constraint on the cooperating parties.¹²⁷

¹²⁶ See, e.g., Case IV/D-2/34.780, *Virgin/British Airways*, Commission decision of 14 July 1999, 2000 O.J. (L/30) 1; dominance at 39.7% of the market, upheld on appeal Case T-219/99, *British Airways v Commission*, 2003 E.C.R. II-5917.

¹²⁷ Horizontal Cooperation Guidelines, *supra* note 94, ¶¶ 137-139.

Third, we should assess entry barriers and potential market entry to look at potential competition. Even if there is no viable *existing* competition, the threat of new entry into the existing technology market could competitively constrain the cooperating parties. For entry to be sufficiently constraining on the parties' behavior, the parties must show it to be likely and timely to pose a viable competitive threat to the parties.¹²⁸

Whether the possibility and likelihood of new entry can constrain the parties will depend on factors such as: (1) whether the cost of entry includes sunk costs or the need to overcome a difficult regulatory framework, like burdensome certification; (2) minimum efficient scale within the industry—if this is large, entry is more costly and risky; (3) whether there has been past entry on a significant scale or not; and (4) whether potential entrants have a realistic chance of competing effectively with the incumbent, (as an illustration, whether they have access to at least as cost efficient technologies).¹²⁹ If we extrapolate about the nature of TAI development based on the development of the most advanced AI today, it is probably not likely to be a market where potential

¹²⁸ *Id.*, ¶ 45-47.

¹²⁹ Article 101(3) Guidelines, *supra* note 56, ¶ 115.

competition is realistic or likely. This is because entry barriers and incumbent advantage are likely large.¹³⁰

To conclude our assessment under Article 101(1), in the innovation market, the R&D cooperation may breach Article 101(1) where the competitor and Leader are cooperating at a late stage of development where the competitor is relatively close to developing TAI independently, and other competitors are no longer viably “in the race.” In an existing technology market, issues may also arise where the competitor and Leader hold a strong position and face little other credible competition or potential competition. This seems to be a possible or even likely scenario given the potential high entry barriers and large incumbent advantage in a TAI development market. If the R&D cooperation restricts competition in the ways outlined above, it may still be exempt if it satisfies the conditions under Article 101(3).

Exception under Article 101(3)

The Article 101(1) prohibition will not apply to the R&D cooperation if the four cumulative conditions under Article 101(3) can be satisfied. The burden of proof rests on the parties seeking to rely on the exemption. We consider each condition in turn below.

¹³⁰ The only exception may be if there are large companies in adjacent markets that have the required resources and related technological acumen that they can leverage into the TAI development market. Even if such a company is technically capable of entering however, it is unclear whether it would be commercially attractive *for it* to do so, given the probable strength/dominance of the cooperating parties in the TAI development market.

The Commission recognizes that many R&D agreements can bring about efficiency gains by leading to “improved or new products and technologies being developed and marketed more rapidly than otherwise,” or bringing about cost reductions.¹³¹ The efficiency gain in this instance, compared to a counterfactual without the R&D cooperation, would be the development of an improved, more rigorously tested version of TAI that is more ethical or safer, to the benefit of consumers.¹³² The R&D agreement could also achieve a reduction in cost, where this cost relates to the cost of consumer harm or litigation as a result of the harmful race dynamic. The challenge, though, is that competition law requires the efficiency gains to be clearly substantiated and shown to flow directly from the R&D cooperation. This is a relatively high evidential standard and includes verifying the direct causal link between the agreement and the efficiencies.

The potential difficulty in relation to the Assist Clause is that the causal link between the R&D cooperation and the resulting efficiencies is more indirect. In other words, can we actually show and quantify that cooperation will lead to better safety outcomes? This may be so if we are combining complementary resources,

¹³¹ Horizontal Cooperation Guidelines, *supra* note 94, ¶ 141

¹³² The following analogous example is given in *id.*, ¶ 149. Two engineering companies pool their R&D efforts to improve the production of an environmentally friendly car component that means that vehicles would consume less fuel and therefore emit less CO₂. In this scenario, the Commission recognizes that the restriction of competition is likely to be outweighed by benefits from consumers from a lower consumption of fuel, thereby satisfying Article 101(3).

skills, or expertise, but may be harder to argue that better safety outcomes would arise because of a reduction in “harmful race dynamics.” This relies on a conjecture that race dynamics are genuinely harmful, whereas competition law generally views “race dynamics” as competition that typically leads to better outcomes for consumers, in the form also of better quality, more innovative (and in this case, safer) products.¹³³ We would therefore need to demonstrate with evidence that race dynamics will create worse AI safety outcomes. One alternative is to base an Article 101(3) argument on more traditional economic efficiencies, such as cost-efficiencies that the R&D cooperation implementing the Assist Clause would incidentally achieve. Cost efficiencies could also arise from developing most cost-efficient production technologies and methods, economies of scale, or combining two existing technologies that have complementary strengths that may reduce production cost or lead to production of a higher quality product.¹³⁴

Any restrictions on competition from the R&D cooperation should not go beyond what is necessary to achieve the efficiency gains. In other words, if there were a realistic alternative¹³⁵ that

¹³³ Note that cooperation to achieve a desirable goal, such as environmental protection, can later stray into an anti-competitive cartel (*see e.g.* Case COMP/39579 Consumer Detergents, Commission decision of 13 April 2011, 2011 O.J. (C 193) 14). This is always a risk when companies work together, even if with good intentions at the outset.

¹³⁴ Article 101(3) Guidelines, *supra* note 56, ¶¶ 64- 68.

¹³⁵ There must be no other economically practicable and less restrictive means of achieving the efficiencies (*id.*, ¶ 75).

would achieve the same gains in safety without reducing the competition on the TAI development market, for instance stepping up monitoring or auditing efforts rather than going as far as the R&D cooperation agreement, then the Assist Clause would not satisfy 101(3). The parties should explain, for example, why they could not have achieved the same efficiencies by acting alone.¹³⁶

Efficiency gains must also be passed onto customers and must outweigh the anti-competitive effects – that is to say, the improved safety must be demonstrated to outweigh any negative consumer effects, like higher prices or lower output, as a result of the loss of competitive dynamics.¹³⁷ A further point to keep in mind is that to satisfy the “fair share” threshold, what counts is the “overall impact on consumers of the products within the relevant market.”¹³⁸ This means that the net effect of a restrictive agreement on the consumers who are subject to the restriction, here users of the eventual AI technology,¹³⁹ must be positive. Whether efficiencies will be passed on to consumers will depend on whether there is enough residual competition on the market and is closely linked to condition 4 discussed below.¹⁴⁰

¹³⁶ *Id.*, ¶ 76

¹³⁷ *Id.*, ¶ 85.

¹³⁸ *Id.*, ¶ 87.

¹³⁹ Furthermore, the term ‘consumers’ includes all users of the goods or services concerned, whether undertakings or private individuals, and at whatever stage of the supply chain (*id.*, ¶ 84).

¹⁴⁰ Article 101(3) Guidelines, *supra* note 56, ¶¶ 95-97.

Finally, the R&D agreement cannot effectively eliminate competition in a substantial part of the products or technologies in question. We should carry out this assessment on both the existing market and the innovation market, and in a similar way to the competitive assessment under Article 101(1).¹⁴¹ The difference is the strictness of the test. Under Article 101(1), the assessment tests whether the agreement appreciably reduces competition compared to the counterfactual. Under this Article 101(3) condition, the test is whether the agreement eliminates, not just reduces, effective competition. The central question, as with the analysis above, is whether there will remain enough other companies competing to develop TAI such that this condition is satisfied. In a way, this leads to a direct trade-off or catch-22 scenario. The greater the residual competition, the greater the likelihood that the R&D agreement will be legally compliant, but the less effective it will be in pursuing the AI governance objective to reduce harmful race dynamics.

On the other hand, if the R&D agreement eliminates all competition for TAI, it will be successful in satisfying its objective from an AI governance perspective but might fall foul of competition law. The better way to think about this conundrum is to seek to significantly dampen race dynamics but leave residual competition such that the cooperation strategy can rely on 101(3).

¹⁴¹ See *supra* Section III.A (Analysis of R&D Cooperation under Article 101(1)).

In any case, this might be the outcome in practice as it is likely that enough companies will not comply with the Assist Clause and continue to compete, but they will need to be viable competitors to the cooperating AI companies.

Category 2: Non-collaborative implementation of Assist Clause

It may be possible to implement the Assist Clause without any collaboration between the competitor and Leader. This could include shifting the R&D from TAI development to another area such as safety or applied research, slowing down TAI development, or disbanding R&D researchers. If these actions are truly unilateral, the implementation will fall outside of Article 101(1). However, even if actions are not *prima facie* or designed to be collaborative, they could still breach Article 101(1) if they involve any agreement, understanding or exchange of information at the point of implementation, including any cooperation with the Leader with respect to the implementation or any exchange of information with the Leader about timing or method. It is therefore important to think critically about whether actions that seem or purport to be unilateral really are in practice.

Assist Clause as a “single and continuous infringement”

If the Assist Clause itself and any subsequent implementation both breach Article 101(1), the EC could treat those breaches as part of a “single and continuous infringement” (SCI).

Two or more agreements or concerted practices may be characterised as a single and continuous infringement of Article 101(1) where three cumulative conditions are met: (1) the agreements or concerted practices share an overall plan pursuing a common objective; (2) each undertaking intends to contribute by its own conduct to the common objective pursued by all participants; and (3) either each undertaking is aware of the offending conduct of the other participants in pursuit of the same objective, or they could have reasonably foreseen that offending conduct would occur and was prepared to take the risk that it would.¹⁴² The EC has confirmed that the fact that the number and identity of participants in the cartel changes over time does not preclude the existence of an SCI.¹⁴³ Each infringing participant may be a party to an SCI for the period during which the regulator can prove that it participated in the infringement.¹⁴⁴

Actions making up an SCI must form part of an overall plan pursuing a common objective. Relevant criteria for assessing

¹⁴² Joined Cases T-204/08 and T-212/08 *Team Relocations and Others v Commission*, 2011 E.C.R. II-3569, ¶ 37. The General Court considered the caselaw and concluded that “three conditions must be met in order to establish participation in a single and continuous infringement: namely the existence of an overall plan pursuing a common objective, the intentional contribution of the undertaking to that plan, and its awareness (proved or presumed) of the offending conduct of the other participants.”

¹⁴³ Case T-377/06 *Comap v Commission* (‘Fittings’) [2011] ECR II_1115, EU:T:2011:108, ¶¶ 85-86.

¹⁴⁴ Concept was first applied in Case IV/31.149, *Polypropylene*, Commission Decision of 23 April 1986, 1986 O.J. (L 230) 1, ¶ 81 *et seq.* and in many cases subsequently.

whether there is an SCI pursuing an overall plan include: (1) identical nature of the products concerned; (2) identical nature of the objectives of the practice at issue; (3) whether the same undertakings participated in the practices at issue (though this is not necessary to find SCI); (4) the identical nature of the rules for implementing those practices; and (5) whether the geographic scope of the practices at issue is the same. Note, however, that this checklist is not fixed or exhaustive, nor is every element necessarily mandatory in every case.

Several consequences flow from a finding of SCI. First, an undertaking that has participated in a “single and complex infringement” as a whole can be liable for the conduct of other undertakings that contributes to the same infringement throughout the period of its participation. Second, the SCI may also allow the EC to penalize conduct that would otherwise have been time-barred. There is a five-year limit for the imposition of a penalty that runs from the end of a period of the SCI. Third, an SCI may only attract one fine, but the level of the fine may be higher to reflect the wider scope or longer duration of a SCI.

In some instances, an SCI finding may turn out to be beneficial for an undertaking because it is liable for a single fine rather than one fine per infringement.¹⁴⁵ In practice, however, the

¹⁴⁵ See Bellamy & Child, *European Union Law Of Competition* 135 (2018).

finding of an SCI is generally detrimental because it is more likely to result in the imposition of a significantly higher fine and wider exposure to follow-on private damage claims.¹⁴⁶

One way of mitigating the risk of breaching competition law with an SCI is to leave each undertaking free to implement the practices in any way they wish, and to avoid coordinating, discussing, or disclosing the rules for such implementation. This may be more feasible than trying to avoid coordinating on the plan or objective, as the latter may be important for trust-building and making AI companies aware of others who have entered into the same obligation.

Conclusion & Mitigating Steps

The Assist Clause may breach Article 101(1) in two ways. First, if there is some form of understanding between AI companies that they should each enter into an Assist Clause agreement, that could constitute an anti-competitive agreement in breach of Article 101(1) because it seeks to reduce competition between competing AI companies, regardless of if or when one of the companies implements it. Second, any subsequent cooperation or agreement between competitors to implement the Assist Clause (such as an agreement to cooperate on R&D) could also breach Article 101(1). The implementation could form a standalone breach. Alternatively,

¹⁴⁶ *Id.*

both entering an Assist Clause agreement and its subsequent implementation could form part of one single and continuous infringement.

The potential anti-competitive effects of implementing the Assist Clause will be particularly problematic where the parties are close competitors and the rest of their competitors are relatively weak. Such a scenario would increase the risk that an R&D cooperation agreement implementing the Assist Clause could breach Article 101(1). Furthermore, the R&D Block Exemption will be challenging to apply if the Leader exceeds the market share threshold, for example because of a strong position in existing and related technologies to TAI. It is possible that R&D cooperation could be exempt under Article 101(3), but it may be challenging to verify that the cooperation will result in efficiencies due to increased safety from reduced race dynamics.

There are a few ways to mitigate the risks. First, companies would not fall within the scope of Article 101(1) in the first place if they acted unilaterally in both committing to the Assist Clause and implementing it. Each AI company should unilaterally decide to be bound by an Assist Clause. For the same reason, companies should, if possible, opt for non-collaborative actions that are truly unilateral to implement it, including slowing down R&D, shifting R&D to other areas, or putting employees on leave. In addition, companies

should not contact competitors to persuade them to enter into or implement the Assist Clause, or otherwise seek to influence or reach an understanding with competitors to that effect.

Second, if unilateral action is not possible, companies should seek to implement the Assist Clause in a way that does not restrict competition. This contradicts the Assist Clause's main objective, but all other things being equal, cooperation should be widespread enough across the TAI development industry to dampen harmful race dynamics, but still leave some sufficient competition to satisfy competition law. It is also better to cooperate at earlier stages of TAI development. The closer the competitor is to being able to develop TAI independently, the greater the loss of competition.

Third, amongst the different ways to implement the Assist Clause, companies can significantly mitigate the risks if cooperation between a competitor and the leading TAI developer does not lead to the loss of the competitor's AI development activities. In other words, instead of shutting down the competitor's AI development or getting rid of its staff, it is preferable for the development activities to shift to another, related area of AI development, like AI safety.

Finally, it is best if the R&D cooperation combines complementary research talent and resources and can achieve economies of scale and scope. Both can generate cost savings or productive efficiencies. These efficiencies are easier to prove and

quantify compared to safety gains and will help to establish a case for exemption under Article 101(3).

B. The Windfall Clause

Key Recommendations

The Windfall Clause does not immediately raise competition law concerns because its proponents do not envisage it to be an agreement between competitors. However, concerns may arise where AI companies want assurance that others will similarly volunteer to be bound by the Windfall Clause in order to ensure a level playing field. Such an “agreement to agree” could infringe competition law to the extent that it disincentivizes companies from generating more profit or expanding sales to avoid triggering the clause. As with the Assist Clause, a company can breach competition law as soon as it enters into the “agreement to agree,” regardless of if or when the Windfall Clause is triggered.

Overall, there seems to be a low risk that the “agreement to agree” will infringe Article 101(1) on a “by effects” analysis because of its potential output-disincentivizing effect. To further mitigate this risk, AI companies should avoid entering into an “agreement to agree” to the Windfall Clause, if possible. This helps to place the conduct outside the scope of Article 101(1) altogether. AI companies could also amend their company objectives, as set out in the corporate articles of association, to the effect that they will

comply in good faith with the Windfall Clause and not seek to circumvent it.

Analysis

The Windfall Clause is one proposal by AI governance researchers to tackle the inequality that may result from development of TAI and its potential to concentrate wealth in the hands of those firms that develop it. The Windfall Clause refers to an agreement entered into by an AI company to redistribute a part of its wealth if and when its earnings exceed a certain threshold such as \$1 trillion in profit, or earnings exceeding 1% of the world's total economic output.¹⁴⁷

The first question is whether the Windfall Clause is a restrictive agreement under Article 101(1). One can structure the Windfall Clause in a way that does not constitute an agreement under Article 101(1). For example, it may be in a private contractual agreement between the AI company and a trust that must collect and distribute the proceeds of the Windfall Clause on behalf of humanity. The obligation to pay would then be between the trust and the AI company. That agreement would not be a qualifying agreement under Article 101(1) so long as the trust is not an

¹⁴⁷ Cullen o'Keefe et al, *The Windfall Clause: Distributing the Benefits of AI*, CENTRE FOR THE GOVERNANCE OF AI RESEARCH REPORT (2020), <https://www.fhi.ox.ac.uk/wp-content/uploads/Windfall-Clause-Report.pdf> ("Windfall Clause Report").

“undertaking,” that is an entity that is active in offering goods or services on a given market.¹⁴⁸

However, firms may only wish to agree to the Windfall Clause if their competitors also do so, in order to ensure a “level playing field.”¹⁴⁹ In this scenario, competing AI firms may wish to enter into a side agreement, with each respectively committing to enter into the Windfall Clause to ensure that no one firm is commercially disadvantaged. This “agreement to agree” to the Windfall Clause could be prohibited under Article 101(1) even if it is not a formal written agreement, but rather some form of informal understanding. However, even if there is an “agreement to agree,” the question is whether the agreement is restrictive of competition. The Clause does not seem to be restrictive of competition per se—the AI companies are just agreeing to give away some of their revenue. However, competition law risks arise where the “agreement to agree” results in reduced competition between companies for fear of triggering the Windfall Clause.¹⁵⁰

¹⁴⁸ See *supra* Section II.A (Article 101(1) Prohibition).

¹⁴⁹ This possibility is considered in the Windfall Clause Report, *supra* note 147, at 25.

¹⁵⁰ Other possible concerns may arise from depending on how the Windfall Clause is structured. For example, there could be issues arising from cross-shareholdings if (1) the Windfall Clause is structured as a superjunior stock grant, and (2) a common fund holds that stock. This could give rise to problems under Article 101(1) if the common fund holds that ‘Windfall Clause’ stock on behalf of multiple AI companies, and that cross-shareholding gives the common fund the ability to coordinate the conduct of those AI companies e.g. because its shares give it a board vote for example in each company. However, this is easily addressed by giving the fund only very clearly defined powers e.g. to distribute the Windfall Clause funds under specific triggering conditions. Or, separate

To some extent, any output-disincentivizing effect will depend on how the Windfall Clause is structured. If a company uses a tiered approach such that only the amount in excess of the \$1 trillion is subject to a percentage levy, and that itself is on a stepped basis, then the disincentivizing effect is less clear than taking a percentage levy of all revenue. It seems that proponents of the Windfall Clause currently prefer a tiered system,¹⁵¹ which also reduces risks from a competition law perspective.

However, even if there is a tiered approach, there may still be disincentives to triggering the Windfall Clause at all. The company may wish to avoid attention to the profitability of the firm, leading to potential additional political and regulatory scrutiny, including under competition law.¹⁵² Another disincentive may come from the signal that a triggering of the Windfall Clause may send to investors. It may alert them that the company must redistribute portion of future earnings and may therefore encourage a flight of investment to competitors that are not bound by the Windfall Clause.

recipient funds can be used for each AI company that is signatory to the Windfall Clause. We do not propose to discuss this potential concern further in this paper.

¹⁵¹ Windfall Clause Report, *supra* note 147, at 10.

¹⁵² For further exploration of these issues, *see id.*, at 25-26; Belfield, *supra* note 16.

Infringement of Article 101(1)

For the reasons discussed above, the Windfall Clause could lead to companies restricting sales or output to avoid triggering the Windfall Clause. In this way, the agreement could constitute an indirect restriction on output in breach of Article 101(1).¹⁵³ The Windfall Clause is unlikely to be an infringement “by object.” It is unlikely that the potential disincentivizing effects from the “agreement to agree” will reveal a sufficient degree of harm to competition such that there is no need to examine its effects.¹⁵⁴ Instead, it seems more likely that the agreement to agree is not a “by object” infringement, and therefore we should fully examine its *effects* on competition to determine whether it is a breach of Article 101(1).

¹⁵³ Competition law looks both at direct and indirect restrictions on competition (the text of Article 101(1) prohibits, *inter alia*, all agreements which “*directly or indirectly fix purchase or selling prices or any other trading conditions*”). An indirect restriction is one that is *de facto* capable of restricting competition on the market. For example, Vertical Guidelines, *supra* note 53, at ¶ 48 explains that resale price maintenance (RPM), that is agreements or concerted practices having as their direct or indirect object the establishment of a fixed or minimum resale price or a fixed or minimum price level to be observed by the buyer gives rise to the presumption that the agreement restricts competition and thus falls within Article 101(1). RPM that are set out in contractual provisions or concerted practices that directly establish the resale price constitute ‘clear-cut’ and direct restrictions of competition. The Vertical Guidelines go on to explain that RPM can also be achieved “through indirect means.” (¶ 48) Examples of these indirect, *de facto* RPM mechanisms include “an agreement fixing the distribution margin, fixing the maximum level of discount the distributor can grant from a prescribed price level, making the grant of rebates or reimbursement of promotional costs by the supplier subject to the observance of a given price level, linking the prescribed resale price to the resale prices of competitors, threats, intimidation, warnings, penalties, delay or suspension of deliveries or contract terminations in relation to observance of a given price level.” (¶ 48) These indirect means incentivise or induce the buyer to comply with the RPM, even if they are not directly obliged to do, and can also breach the Article 101(1) prohibition.

¹⁵⁴ Case C-67/13 P, *Groupement des Cartes Bancaires v Commission*, *supra* note 51, ¶ 57.

In the present case, the disincentivizing effect of the “agreement to agree” could mean that the undertakings compete less vigorously on both output and innovation compared to the counterfactual.¹⁵⁵ To constitute a breach of Article 101(1), this adverse effect must be both likely—that is, expected with a reasonable degree of probability—and appreciable—that is, not insignificant.¹⁵⁶

However, any “agreement to agree” that includes the Big Tech companies is likely to lead to appreciable anticipated effects, given their size on the market.¹⁵⁷ Restrictive effects on competition, that is to say whether the restriction is sufficiently likely, will largely depend on how the Windfall Clause is structured. In this regard, it is advisable to structure the Clause in a way that mitigates the probability or magnitude of any disincentivizing effect as much as possible.

¹⁵⁵ Article 101(3) Guidelines, *supra* note 56, ¶ 24.

¹⁵⁶ This requirement of “appreciability” is given expression in the Commission Notice on agreements of minor importance which do not appreciably restrict competition under Article 81(1) of the Treaty establishing the European Community (*de minimis*), 2001 O.J. (C 368) 13 (“*De Minimis* Notice”). The notice states that agreements will not appreciably restrict competition if the aggregate market share of the parties to the agreement does not exceed 10%, where the agreement is made between actual or potential competitors (¶ 8). However, the *de minimis* doctrine does not apply to “by object” restrictions (¶ 2). Further, the Horizontal Cooperation Guidelines state that “restrictive effects on competition within the relevant market are likely to occur where it can be expected with a reasonable degree of probability that, due to the agreement, the parties would be able to profitably raise prices or reduce output, product quality, product variety or innovation. This will depend on several factors such as the nature and content of the agreement, the extent to which the parties individually or jointly have or obtain some degree of market power, and the extent to which the agreement contributes to the creation, maintenance or strengthening of that market power or allows the parties to exploit such market power” (¶ 28).

¹⁵⁷ Horizontal Cooperation Guidelines, *supra* note 94, ¶¶ 39-47.

Overall, however, it may be difficult for a regulator to show to the requisite evidentiary standard that a disincentivising effect is “likely.” The EU courts are clear that the EC must support a “restrictive effects” finding with robust evidence of the effects alleged.¹⁵⁸ This requires some speculation about firms’ behavior and incentives, which may be difficult to back up with convincing evidence. Unless the Windfall Clause is clearly structured in a way that would have a disincentivising effect, there is a low risk that the “agreement to agree” will infringe article 101(1) on a “by effects” analysis.

Conclusion & Mitigating Steps

Overall, there is likely to be a low risk of an “agreement to agree” to the Windfall Clause breaching Article 101(1). This risk will hinge on the extent of the disincentivizing effect discussed above. The best way to mitigate the risk is to do away with the “agreement to agree” altogether. Further research should explore whether this is commercially viable: whether companies will agree to the Windfall Clause even if they cannot be sure that their competitors will do the same.

¹⁵⁸ For example, in *European Night Services v Commission* the General Court rejected the EC’s finding that the establishment of a JV could restrict actual or potential competition between its parents. The Court considered this to be “a hypothesis unsupported by any evidence or any analysis of the structure of the relevant market from which it might be concluded that it represented a real, concrete possibility.” (Cases T-374/94 etc, *European Night Services v Commission*, 1998 E.C.R. II-3141).

Other ways to mitigate could involve reducing the disincentivising effect of the Windfall Clause itself. As well as taking a tiered approach to the “taxing” of any windfall, other mitigation steps could include incorporating a good faith obligation to implement the Windfall Clause into a company’s articles of association. Any attempts to circumvent it, or avoid triggering the Clause, would then be a breach of the company directors’ duties. The right publicity around the Windfall Clause may also be important, so that, for example, one could see the triggering of the Windfall Clause as a positive fulfilment of a company’s corporate social responsibility (CSR).¹⁵⁹ The objective is that such publicity might lead the Windfall Clause to encourage investment in the firm in question, rather than investors seeing it as a negative funnelling away of profits away from potential payout in dividends.¹⁶⁰

IV. LEGAL ANALYSIS OF PROPOSED FORMS OF COOPERATION: GENERAL TRUST BUILDING

In this section, we analyze cooperation strategies to encourage safe and responsible AI development include mutual or

¹⁵⁹ This could be plausible especially in light of current growing interest around sustainable, responsible and impact investing (SRI), and so may in fact encourage investment from certain, socially impactful investors. For example, according to the Global Impact Investing Network (GIIN), the global impact investing market has grown to more than US\$500 billion by 2019, more than doubling from an estimated US\$ 228 billion in 2018. See H. Liang and L. Renneboog, *Corporate Social Responsibility and Sustainable Finance: A Review of the Literature (September 24, 2020)*, European Corporate Governance Institute – Finance Working Paper No. 701/2020, forthcoming in the Oxford Research Encyclopedia of Economics and Finance.

¹⁶⁰ Note in a draft companion piece we assess AI governance proposals that might intersect with abuse of dominance, public procurement and/or state aid concerns. See *supra* note 35.

third-party safety monitoring and standard setting. Companies should ideally implement these strategies as soon as possible, as they address AI risks from both a long-term and near-term perspective. We survey three types of cooperation: agreements, information sharing, and standard setting. As before, for each type of cooperation we explain what it is, why we would want it from an AI governance perspective, what competition concerns it might raise, and how to address those concerns.

A. Agreements Between Competitors

Cooperation between competitors can take the form of an “agreement” that is specified and formalized in a legal agreement, or written declaration, a verbal agreement, an “agreement to agree” or any other form of understanding. Agreements can raise competition concerns under Article 101(1) if they are capable of restricting competition. The Assist Clause and the Windfall Clause are both agreements, and we have assessed above their compatibility with Article 101(1). In this section, we will examine a third type: agreements on secure enclaves.

*I. Agreements on Secure Enclaves*Key Recommendations

Competition law will generally view an agreement between AI companies and hardware companies that do not compete with each other as pro-competitive. Any agreements that competing AI companies enter into on secure enclaves, however, are likely to constitute a “horizontal standards” agreement between competitors.¹⁶¹

Analysis

One proposal within AI governance is that AI and hardware companies work together to explore the feasibility of developing secure enclaves for application in AI accelerators, or otherwise devise best practices on how to use secure enclaves, otherwise known as Trusted Execution Environments, in machine learning contexts.¹⁶² A secure enclave is a set of software and hardware features that together provide an isolated execution environment that enables a set of strong guarantees for applications running inside the enclave. Secure enclaves provide guarantees of security, provided that their underlying assumptions cannot be broken. These mechanisms help to focus defensive efforts and assure users that relatively extreme measures would be required to subvert their trust.

¹⁶¹ Assessment of horizontal standardization agreements are further discussed in *infra* Section IV.C (Standard-setting).

¹⁶² Brundage, Avin, Wang *et al.*, *supra* note 4.

If the collaboration is between AI and hardware companies, it is possible that they are not competitors. This is because AI software and hardware are complementary rather than substitute products. Competition law generally deems collaboration between suppliers of complementary products to be pro-competitive and efficiency-enhancing (so-called “vertical agreements”).¹⁶³

It may be that after AI companies work with hardware companies to develop best practices on secure enclaves, AI companies seek to agree a standard or “best practice” between each other. This may constitute a horizontal agreement on safety or technical standards between competing AI companies, in contrast to the vertical agreement between AI and hardware companies.¹⁶⁴ B. Information exchange

B. Information Exchange

A variety of AI risk mitigation strategies involve monitoring AI labs and sharing information (such as incident sharing or compute costs). Although these proposals pursue legitimate aims, competition law risks can arise if this monitoring and information sharing includes the exchange of commercially sensitive

¹⁶³ See, e.g., Horizontal Cooperation Guidelines, *supra* note 94, ¶ 50. For further discussion of vertical agreements, see *supra* Section II.A (Article 101(1) Prohibition).

¹⁶⁴ Horizontal Cooperation Guidelines, *supra* note 94, ¶ 55. See *infra* Section IV.C (Standard-setting) for further assessment of standardization agreements.

information between competitors, whether directly or indirectly through a third party.

Generally, any exchange between competitors of information that reduces strategic uncertainty in the market can constitute an unlawful information exchange under Article 101(1).¹⁶⁵ This can lead to two key anti-competitive effects: first, it may enable undertakings to predict each other's behavior and thereby coordinate their behavior on the market; and second, undertakings may use it to support a cartel, for example, to spot deviation and ensure compliance with a price-fixing agreement.

Such “strategic” information can include (but is not limited to) the following types of non-public information: (1) prices, including wholesale or retail prices, pricing strategy, discounts, margins, costs; (2) sales and volumes, including sales volumes, revenues, stock levels, market share, production capacity; (3) R&D, including details of a company’s technology, R&D programmes and their results, investments; (4) product strategy, including product development or marketing plans; commercial strategy information (geographic growth and business expansion or contraction plans); and (5) customers and bidding strategy, including customer lists, individual suppliers, terms and conditions of supply, bid amounts,

¹⁶⁵ Horizontal Cooperation Guidelines, *supra* note 94, ¶ 86.

terms, and the decision on whether or not to respond to a particular tender.

In this section, we consider forms of information exchange or auditing directly between competitors, information exchange with a third party, and information exchange with “bounty hunters.”

1. Mutual Monitoring

Key Recommendations

Mutual monitoring by one AI company of another competing AI company may raise competition law concerns where it involves the exchange of commercially sensitive information such as R&D plans, product roadmaps, non-public insight into how the technology or AI system works. Using a third party rather than a competitor to carry out this monitoring can significantly reduce the competition law risk. This could be a research institute or an association such as PAI.

Analysis

One AI governance proposal is for AI companies to monitor or audit other AI companies’ development.¹⁶⁶ The benefits of using other AI companies to monitor in this way is that they may be equipped with the necessary scarce resources and expertise to be able to carry out this type of monitoring to the requisite standard,

¹⁶⁶ Brundage, Avin, Wang *et al.*, *supra* note 4.

and it may assure them that others are not “cutting corners” on responsible development.

Mutual monitoring seeks to pursue a legitimate and pro-competitive purpose; it may require sharing of non-public information about the AI company’s R&D, technology, or algorithms. As the Horizontal Cooperation Guidelines explain, if companies compete on R&D then the exchange of technology data is most likely to be risky from a competition law perspective. The risk is that in monitoring the safety or ethics of an AI-related technology, a competitor will also gain insight into its competitor’s AI technology, such as the proprietary algorithm. This potentially breaches Article 101(1).

As well as the nature of the information exchanged under the monitoring, the strategic usefulness (and therefore level of competition law risk) will depend on several other factors. First, it is important whether the information is public or non-public. The exchange of genuinely public information is a far lower risk than exchanging confidential information. However, presumably the information used for mutual monitoring is more likely to be the latter because AI companies will wish to keep information on their technology confidential to protect their IP. This could increase the competition law risk.

Second, aggregated data will be lower risk than individualized data. However, it may be difficult to aggregate the monitoring information, as the concern is development within an individual company rather than in the industry more generally. Third, the older the data, the lower the risk.¹⁶⁷ Although the exchange of present or past safety results, for example, is lower risk than exchanging information on future R&D development, risks may still arise if it remains indicative of competitors' future conduct. The data will presumably be relatively recent in order to be useful for monitoring purposes. This may increase the competition law risk. Fourth, the exchange must affect a sufficiently large part of the relevant market to be capable of having a restrictive effect on competition (the threshold is likely to be 10% combined market share between the companies exchanging the information).¹⁶⁸ As we are most interested in monitoring the leading AI development companies, they are also more likely to have a significant share of the relevant AI development market.¹⁶⁹

The above analysis explains when the information exchanged as part of mutual monitoring may be strategically significant: when it gives the auditing company insight into the

¹⁶⁷ Generally, the EC will regard information that is more than one year old as historic, although in a fast-moving market such as AI development it is likely to age far quicker. *See* Horizontal Cooperation Guidelines, *supra* note 94, ¶ 90 n.67.

¹⁶⁸ De Minimis Notice, *supra* note 156.

¹⁶⁹ *See supra* Section II (Background: The Prohibition on Anti-competitive Agreements under Article 101(1)).

likely future behavior or market strategy of a competitor. Information exchanges are especially capable of restricting competition on the market where they are non-public, non-aggregated, relatively recent, and cover a large part of the market. Nevertheless, competition law recognizes that some exchanges of commercially sensitive information can be efficiency-enhancing and pro-competitive. The analysis below will consider the extent to which the efficiencies, such as the safety enhancements arising from the monitoring mechanism, can satisfy the four requirements of Article 101(3).

The exchange of safety development information would achieve a pro-competitive benefit under the first conditions of Article 101(3) because it should improve safety.¹⁷⁰ However, another important criterion to satisfy under Article 101(3) is indispensability, that is that the restrictions on competition do not go beyond what is necessary to achieve the efficiency gains.¹⁷¹

In this scenario, it seems that the mutual monitoring could feasibly be done by a third party, such as a pure research institute, that is not an active “undertaking” that either does not provide goods

¹⁷⁰ Case IV/33.863, *Asahi/Saint-Gobain*, Commission Decision of 16 December 1994, 1994 O.J. (L 354) 87, ¶¶ 24-34 (the Commission found that even though there were competitive risks with regard to a joint venture, it satisfied the conditions of Article 101(3) because the new type of car glass would contribute to driver safety, and the lighter weight would help fuel efficiency. The cooperation between the parties had the effect of reducing the R&D costs for this new car glass and the entry of such products on the market would be accelerated.)

¹⁷¹ Horizontal Cooperation Guidelines, *supra* note 94, ¶ 101. For further discussion of these other forms of mutual monitoring.

or services on the market or is in the AI market. This would be a realistic and less restrictive means of achieving the same objective. Moreover, it is not enough to simply use the third party as a conduit to pass the same commercially sensitive information between the same competing companies, as that could itself be an unlawful cartel or information exchange.¹⁷² If this is done via a trade association, the dissemination of sensitive information by the trade association is treated as an information exchange or concerted practice by the members themselves.¹⁷³ The companies must trust the third party to carry out the monitoring and perhaps to publish results of level of compliance without disclosing commercially sensitive information, unless the third party releases it into the public domain.

To satisfy the other grounds of Article 101(3), companies should pass on the efficiency gains to consumers, and the companies involved in the exchange must not eliminate competition in a substantial part of the relevant market. These two conditions should be satisfied so long as there are some competitors on the market that are outside of the monitoring system.

¹⁷² For example, in 2008 the European Court of First Instance affirmed the Commissions' decision that a Swiss consultancy firm, AC Treuhand, had breached competition law for facilitating an organic peroxide cartel, finding that Treuhand's role, though not active in the market of the relevant product, was that of a passive co-perpetrator and so was caught by Article 101" (Case T-99/04, AC-Treuhand AG v Commission, 2008 E.C.R. II-1501, ¶ 122).

¹⁷³ See, e.g., Case 45/85, *Verband der Sachversicherer v Commission*, 1987 E.C.R. 405, EU:C:1987:34; *id.*, ¶ 105.

Conclusion & Mitigating Steps

The information exchange involved in mutual monitoring risks breaching Article 101(1). Furthermore, it may be difficult to rely on the Article 101(3) exemption. This is because, unless there is a good reason why competitors must mutually monitor each other, it seems that the exchange will not satisfy the “indispensability” limb of Article 101(3), as there are less restrictive alternatives such as using a third party to monitor instead.

Overall, using a third party rather than a competitor to carry out the monitoring can reduce the competition law risk. That third party could be a research institute or an association such as PAI. If that is not feasible, for instance because only competing AI companies have the necessary technology and skills to effectively monitor, the parties should ensure there are contractual and operational safeguards that information shared is only used in those companies for the immediate purpose. In addition, the monitoring should only share information that is strictly necessary for the monitoring purposes, and nothing more. Other ways to mitigate the risk include sharing information that is (as far as is possible) public, historic, aggregated and/or anonymized and ensuring that every company is able to access this initiative (to the extent a company gets a competitive advantage from being monitored or receiving results of other companies’ monitoring results).

2. *Other Forms of Mutual Monitoring*

Other forms of mutual monitoring include shared red teaming or “white-hat” hacking to uncover vulnerabilities; incident sharing; and compute accounting.

These all raise a similar concern to mutual monitoring in the previous section—that monitoring will lead to the sharing of competitively sensitive information between competitors. One can mitigate them in a similar way, most importantly by engaging a third party to carry out the monitoring instead. This is particularly advisable for compute accounting, where the information exchanged is a key component of cost and therefore of high commercial sensitivity and commensurately higher competition law risk.

3. *Communication, especially “Heads-up”*

Key Recommendations

The “heads-up” strategy could raise competition risks to the extent that it involves the exchange of strategically important, forward-looking R&D data. One could also see it as an agreement between competitors to restrict or delay output.

One can significantly mitigate risks if a standard-setting organization (SSO) carries out any pre-publication review based on defined, objective publication norms or criteria.¹⁷⁴

¹⁷⁴ See discussion *infra* Section IV.C.2 (Publication and release norms).

Analysis

Greater communication between competitors can be valuable from an AI governance perspective. This could include regular updates, clear communication channels and people responsible for communication¹⁷⁵, and joint events such as workshops, retreats, conferences, residencies, country visits.

In the OpenAI report on lessons learned from GPT-2, the authors emphasized the importance of communication, recommending:

“Build communication channels across organizations. Research results are often kept private until the associated paper is published. Private results hinder coordination, especially for release; for example, we were largely unable to retrieve statuses of replication efforts. The norm of privacy around unpublished research holds legitimacy, as seen in non-disclosure agreements, but robust communication channels between AI organizations will be needed in the future. For example, prior to first announcing GPT-2, we were unsure whether and how quickly other labs would eventually develop and publish similar systems. Since the impact of an individual publication decision often depends on others’

¹⁷⁵ Brundage, Avin, Wang *et al.*, *supra* note 4, at 9 (“[C]lear identification of roles or offices within organizations who are responsible for maintaining and deepening interorganizational communication.”).

publication decisions, we encourage AI labs to experiment with their approaches to interorganizational communication.”¹⁷⁶

This subsection will focus on a pre-publication or pre-release “heads-up” to competitors about R&D or products. These are intended to give competitors an opportunity to point out reasons why the R&D product should be stalled, for example because of outstanding safety concerns or because it can be exploited by malicious actors.¹⁷⁷ Including others and giving them an opportunity to express concerns can build a culture of scrutiny and accountability amongst AI firms, especially in AI safety.

However, communications and “heads-up” practices may raise competition law risks and are potentially higher risk than the examples of information exchange discussed above for several reasons. First, the “heads-up” could involve sharing future R&D or product releases. Future information is more commercially sensitive than past information because it gives direct insight into a competitor’s future behavior, allowing the competitor to adapt their own commercial strategy accordingly.

Second, the “heads-up” may go beyond exchange of commercially sensitive information and provide an opportunity for competitors to reach an agreement to limit or delay output. Both of

¹⁷⁶ Irene Solaiman et al., *Release Strategies and the Social Impacts of Language Models*, OPENAI REPORT 24 (November 2019), arXiv:1908.09203.

¹⁷⁷ For example, several groups and companies have not released full language models, starting with GPT-2.

these factors could make this a “by object” infringement, which refers to agreements that are by their very nature sufficiently harmful to competition that there is no need to look at its anti-competitive effects.¹⁷⁸ Although a “by object” infringement can still theoretically benefit from the Article 101(3) exemption—indeed, any agreement can theoretically benefit from the exemption¹⁷⁹—it is in practice very difficult for it to do so.¹⁸⁰

In addition, the “heads-up” strategy may struggle to satisfy the conditions under Article 101(3). First, to satisfy the “efficiencies” condition, one could argue that the “heads-up” generates efficiencies by reducing harmful disclosure by limiting disclosure that might facilitate the malicious use of AI by criminals or terrorists. This security improvement may constitute “an improvement in the production or distribution of goods or in technical or economic progress” that outweighs any detrimental effects on competition.¹⁸¹ It is not advisable to argue efficiencies from building interpersonal trust between AI labs, because it is

¹⁷⁸ See, e.g., Case COMP/39188, *Bananas*, Commission decision of 15 October 2008, 2009 O.J. (C189) 12, upheld on appeal, Case T-588/08, *Dole Food v Commission*, EU:T:2013:130, upheld on further appeal, Case C-286/13P, *Dole Food v Commission*, EU:C:2015:184. It is also well-established that agreements between competitors to fix prices, limit output or share markets are restrictive by object: see e.g., Article 101(3) Guidelines, *supra* note 56, ¶ 21 and Case C-209/07, *Competition Authority v Beef Industry Development Society*, 2008 E.C.R. I-8637, EU:C:2008:643 (agreement between competing companies to reduce capacity was held to be a ‘by object’ infringement).

¹⁷⁹ See, e.g., Case C-439/09, *Pierre Fabre Dermo-Cosmetique SAS*, *supra* note 59.

¹⁸⁰ See Horizontal Cooperation Guidelines, *supra* note 94, ¶ 46.

¹⁸¹ Cases 56/64 and 58/64, *Consten and Grundig v Commission*, 1966 E.C.R. 299, at 348; Case T-65/98 *Van den Bergh Foods Ltd v Commission*, 2003 E.C.R. II-4563, ¶ 139.

rather speculative and is a less-clear cut efficiency compared to security.

However, as mentioned above, a narrow view of Article 101(3) only permits agreements that would bring about improvements in economic efficiency. The Article 101(3) Guidelines clearly support this view.¹⁸² Even assuming that a broader view of efficiencies is possible, it is also unclear whether we would be able to prove and quantify the alleged security gains with “convincing arguments and evidence” to fulfil the standards of Article 101(3) with evidence like expert reports, or economic analysis.¹⁸³ It may be particularly difficult to show that a purported gain in security has a sufficient causal link to the “heads-up” system such that the reduction in potential malicious use of AI is sufficiently likely to arise from it.¹⁸⁴

Second, even if the “efficiencies” condition is satisfied, it may be difficult to establish that there are no less restrictive

¹⁸² Article 101(3) Guidelines, *supra* note 56.

¹⁸³ *See, e.g.*, Case COMP/34579 MasterCard, Commission decision of 19 December 2007, ¶¶ 686, 690, 732. The decision explains:

A claim that an interchange fee mechanism creates efficiencies...must be founded on a detailed, robust and compelling analysis that relies on its assumption and deductions on empirical data and facts...general references to economic theory [that] do not go beyond paraphrasing (partly inaccurately) some general conclusions that might be drawn from economic models, but do not specify the models used, the assumptions relied on and the facts and data that support the analysis...does not present a detailed and robust economic and empirical analysis.

¹⁸⁴ *See, e.g.*, Cases C-501/06P, etc, GlaxoSmithKline Services v Commission, 2009 E.C.R. I-9291, EU:C:2009:610, ¶ 94 (the ECJ held that the Commission’s approach may involve determining whether “it seems more likely either that the agreement in question must make it possible to obtain appreciable advantages or that it will not.”).

alternatives to achieve the same aim as the “heads up.” For example, companies could share pre-publication R&D or products plans with a third-party organization that provides an objective view (or even certification) as to whether it satisfies a requisite standard. This seems to be a viable alternative that both achieves the legitimate purpose in a more consistent and effective way while significantly reducing the competition law risk.

Conclusion & Mitigating Steps

Given the breadth of the Article 101(1) prohibition, even a “heads-up” in the form of an informal notification or request for feedback from one AI company to another could raise risks under Article 101(1). These risks could arise if the “heads-up” involves an exchange of commercially sensitive information—for example, forward-looking R&D data—that reduces one company’s uncertainty about its competitor’s commercial strategy. If the “heads-up” strategy provides an opportunity for competing AI labs to agree or reach an understanding on restricting or delaying the launch of a technology, this could also give rise to material competition law risks because it potentially constitutes an agreement between competitors to restrict output.¹⁸⁵

¹⁸⁵ Of course, it is possible that labs may (and in fact do) disagree about the safety implications of some R&D data, even in the context of heads-ups being given. If no agreement or ‘meeting of minds’ occurs on whether or not to release the data, then Article 101(1) is unlikely to apply. However, even if no agreement is reached, the discussions could still involve exchange of commercially sensitive information (e.g., R&D data) which raises separate competition law risk.

Furthermore, it may be challenging to argue that the agreement satisfies Article 101(3). First, the increase in safety may be difficult to evidence and quantify. Second, there are likely to be less restrictive means available to achieve the same aims as the “heads-up” strategy.

A lower risk alternative could be to send the pre-publication material to an SSO such as PAI for review based on pre-defined, objective publication norms or criteria.¹⁸⁶

4. *Seconding Staff*

Key Recommendations

To help build trust between AI companies, AI companies could second their researchers or engineers to competing AI labs.

Several kinds of joint projects between competitors could be valuable from a cooperation perspective, such as technical projects, research efforts and papers, policy projects, projects to serve some social good, or efforts to expand access. This analysis will look at a specific scenario where secondment of staff is part of an R&D cooperation agreement, for example to implement the Assist Clause.

Analysis

The secondment of staff to a competitor necessarily involves the exchange of commercially sensitive information. The staff

¹⁸⁶ Note that this is the model proposed for AI systems in high-risk sectors and application areas in the *EU White Paper on Artificial Intelligence: A European Approach to Excellence and Trust* (Feb. 19, 2020), https://ec.europa.eu/info/sites/info/files/commission-white-paper-artificial-intelligence-feb2020_en.pdf.

member is likely to receive some information about the competitor firm's R&D plans, technology, costs, or capacity. The staff member may also wish to share information about their "home" company's R&D or technology to legitimately assist in the joint R&D project. At the same time, this exchange of commercially sensitive information also risks breaching Article 101(1). The information exchange reduces the strategic uncertainty between the competing AI labs if the information were to be used to inform the labs' overall R&D and commercial strategy. However, where the information exchange is part of a legitimate agreement or transaction, such as an R&D cooperation agreement that has been found to be lawful, it may qualify as an "ancillary restraint" and fall outside the Article 101(1) prohibition altogether.¹⁸⁷

The information exchange will be an ancillary restraint falling outside Article 101(1) if it is necessary and proportionate to the implementation of a main, legitimate agreement. To qualify as an ancillary restraint, the exchange must first be "directly" related to the R&D cooperation agreement. For this condition to be satisfied, the exchange must be subordinate to the implementation of that transaction, and inseparably and evidently linked to it.

¹⁸⁷ We have not assessed this as a standalone information exchange because we assume it will be part of some kind of R&D cooperation agreement. If it is standalone, it should be assessed under normal Article 101(1) principles especially those relating to information exchange- the assessment in section above on auditing/information exchange should apply similarly in this instance.

Second, the exchange must satisfy the requirement of objective necessity, that is one must show that without the information exchange, the R&D cooperation would be *impossible* to carry out.¹⁸⁸ It is not sufficient if it is merely more difficult or less probable.¹⁸⁹ Third, the information exchange must be proportionate in duration and scope to the implementation of the main transaction. This means there should only be the minimum amount of information exchange strictly necessary to achieve the legitimate purpose of the R&D cooperation—that is, to jointly pool engineer resources to boost safety development.¹⁹⁰ To help ensure these conditions are satisfied, it is important to take a number of mitigating steps outlined below.

Conclusion & Mitigating Steps

In theory, the information exchange related to the staff secondment should be a lawful ancillary restraint that falls outside Article 101(1), subject to putting the safeguards above in place. However, in reality, the risk is that there will be “leakage” of information beyond the confines of what is lawful, which is natural when a secondee is mixing with a competitor AI lab’s employees on

¹⁸⁸ An ancillary restraints assessment borrows some concepts from the exemption under Article 101(3) such as ‘necessity’. However, an important distinction is that it does not require the challenging exercise of weighing up pro- and anti-competitive effects. That exercise is reserved exclusively for Article 101(3) (Article 101(3) Guidelines, *supra* note 56, ¶ 31).

¹⁸⁹ *Id.*, ¶ 29; Case T-112/99, *Metropole Television (M6) v Commission*, 2001 E.C.R. II-2459, ¶ 109.

¹⁹⁰ Brundage, Avin, Wang *et al.*, *supra* note 4, ¶ 9.

a day-to-day basis. Therefore, this strategy could give rise to competition law risk.

Several mitigating steps are possible. First, to ensure that the information exchange relates to only what is sufficiently related and objectively necessary, the secondees must be sufficiently “ring fenced,” for example under strict contractual obligations to use the competitor’s commercially sensitive information only for a legitimate purpose of enhancing safety development. They should not be using it for any other purpose, such as to inform the secondee’s company’s own R&D plans unrelated to enhancing safety development.

Similarly, the secondee should disclose their own company’s commercially sensitive information only for the legitimate purpose. There should also be corresponding contractual obligations between the two companies that any information that they receive about each other can only be used for the legitimate safety purpose and no other, and especially not to inform each other’s commercial strategy. Practical safeguards should also bolster contractual ones. For example, the secondee should be given sufficient and targeted competition law training and should be physically located in a team that is focused on, for instance, safety R&D and no other functions. In particular, the secondee should be

separated from teams that are involved in commercial strategy, like sales and marketing teams.

However, even with the most stringent contractual and practical restrictions on information exchange, it may be difficult in practice to guarantee that the secondee never shares any more information than is necessary for the legitimate purpose, and that the companies receiving the competitor information guards against the temptation to use that information to generally inform its commercial strategy. Given the potential risks, it may be advisable to opt for alternative ways to achieve the same purpose. For example, using a third party such as PAI to send its own staff in to carry out quality control may be one option that significantly lowers the competition law risk.¹⁹¹ PAI could then draw up a periodic “trust report” to review how “trustworthy” or reliable they believe AI development is in the different AI labs.

5. *Third-party Auditors*

Key Recommendations

Using a third-party auditor (such as PAI) is a less risky alternative to mutual monitoring between competing AI companies. However, companies should not use the third-party organization as a mere intermediary for the competing firms to exchange the same commercially sensitive information. Instead, the third-party

¹⁹¹ See *infra* Section IV.B (Information Exchange).

organization should carry out the auditing on an independent basis, without disseminating any commercially sensitive information between competitors. Alternatively, the organization could collect commercially sensitive information and then disseminate that to the competing parties in a “sanitized” form: aggregated or anonymized.

Analysis

The process of AI development is often opaque to those outside a given organization and various barriers make it challenging for third parties to verify the claims developers make. As a result, it may not be easy to verify claims about system attributes. AI developers have justifiable concerns about being fully transparent with information concerning their AI development, for example to prevent malicious use. One potential solution is to set up or use a third-party association or standards body that could carry out independent, third-party auditing of AI systems. AI companies could legitimately fund this, so long as it nevertheless operates as a separate and independent body.

As mentioned above, many forms of cooperation and auditing that we consider in this paper are more likely to comply with competition law if a third-party organization rather than another AI company carries them out. However, this depends on two key points. First, any third-party organization will not mitigate competition law risk if simply used as a “cover” or vehicle for collusive conduct or conduit for illegal information exchange

between competitors. However, where an organization such as PAI carries out the third-party audit in a genuinely independent manner, this is likely to be compliant with competition law.

Second, a trade association or cooperative body will run risks under competition law if it restricts the ability of market participants to compete. Therefore, to the extent that a third-party auditing or certifying companies in this way would confer a significant competitive advantage, all AI companies should have the opportunity to do so. If the audit is restricted, for example to PAI members only, that means membership of PAI is likely to be an essential or important pre-conditions to operating on the market. It follows that unreasonable refusal to admit members could restrict competition and infringe Article 101(1). One should therefore design any restrictions on membership rules to fall within the Article 101(3) exemption. To help satisfy the indispensability requirement in Article 101(3), the admission rules must be objective, sufficiently determinate, and capable of uniform and non-discriminatory application.¹⁹²

Conclusion & Mitigating Steps

Generally, as mentioned above, it is less risky to use a third-party organization to carry out some of the envisaged risk mitigation auditing exercises. The exchange of competitively sensitive

¹⁹² Cases T-528/93 etc, *Metropole Television v. Commission*, 1996 E.C.R. II-649, EU:T:1996:99, ¶ 102.

information or other anti-competitive agreements leads to higher competition law risk when competitors carry it out between themselves. They can significantly mitigate risks if auditing is delegated to a third-party organization.

However, this is not a “catch-all” solution. If the competing firms simply use the third-party organization as an intermediary to carry out an anti-competitive cartel or information exchange, then that arrangement will still breach Article 101(1). To reduce competition law risk to a low, manageable level, the third-party organization should: (1) operate as an independent and distinct entity in carrying out the audits, vis-a-vis the AI labs that fund it; (2) make the possibility of audit open to all AI labs (or in the case of PAI, only to PAI members but subject to the criteria above on membership rules); (3) not disseminate any commercially sensitive information that it receives as a result its audit to any other AI labs, except in a “sanitized” format.

6. *Bias and Safety Bounties*

Key Recommendations

Competition risks can arise where the bounty hunter discloses confidential information about its technology or R&D in reporting how its systems exposed the vulnerability. However, on its own, such information may not be strategically significant and may be justified under the Article 101(3) exemption.

Analysis

“Bug bounty” programs, in which organizations commit to giving individuals recognition or compensation for reporting bugs, have become popular in the information security industry as a way to compensate individuals for recognizing and reporting bugs, especially those related to exploits and vulnerabilities. Bug bounties enable individuals to report bugs directly to the institutions affected, rather than immediately exposing the bugs publicly or selling the bugs to others. To earn the bounty, individuals typically need to articulate the scale and severity of the bugs in order to determine appropriate compensation. Researchers have suggested the establishment of a similar system for bias and safety problems around AI systems.¹⁹³

The level of competition risk largely depends on how the programme is structured. The first question is whether there is an exchange of any competitively sensitive information. For example, an individual simply accessing an AI company’s open- source software, discovering its vulnerabilities, and reporting them, could earn the bug bounty. This does not involve any disclosure of competitively sensitive information, therefore would not give rise to any competition law risk. Risks may arise, however, if in reporting the bug the individual or reporting company discloses commercially

¹⁹³ Brundage, Avin, Wang *et al.*, *supra* note 4, 16.

sensitive information about its own systems or processes. However, this can be mitigated to the extent that it shares only as much information as is strictly necessary for the legitimate purpose of enabling better bias or safety development.

Another relevant question is the relationship between the organization providing the bounty, and the “bounty hunter”. If the bounty hunter is not an actual competitor or potential competitor to the relevant AI company then any information exchange is less likely to raise concerns.¹⁹⁴ To be a “potential competitor,” it must be likely that the bounty hunter can within a short period of time undertake the necessary additional investment or necessary switching costs to enter into the relevant market, if there were a sufficient profit incentive to do so.¹⁹⁵

Conclusion & Mitigating Steps

The competition law risks for a bug bounty programme will generally be low as long as: (1) there is no or minimal exchange of commercially sensitive information; (2) that commercially sensitive information is made public; or (3) there is no exchange between actual or potential competitors on the relevant market. However, take a scenario where there is some non-public exchange of

¹⁹⁴ Horizontal Cooperation Guidelines, *supra* note 94, ¶ 10. *See also* discussion on “horizontal and vertical agreements” *infra* Section II (Background: The Prohibition on Anti-Competitive Agreements Under Article 101(1)).

¹⁹⁵ Further, this assessment must be based on realistic grounds, and not just a mere theoretical possibility (*id.*, ¶ 10).

commercially sensitive information between competitors because the bounty hunter has disclosed how its systems exposed the vulnerability and the bounty hunter is a competitor. Even where none of the above three conditions apply, and although the disclosure provides some insight into the technical system of the bounty hunter, it may be justifiable under Article 101(3). This may be possible on the basis that this is not highly strategic information, and the hunter has exchanged the information that is necessary to achieve the objective.

C. Standard-Setting

Competition law usually recognizes standardization agreements, defined as technical or quality requirements with which products, production processes or methods should comply, as pro-competitive.¹⁹⁶ These agreements can achieve positive economic effects such as promoting innovation, enhancing interoperability, and maintaining and enhancing quality.¹⁹⁷ However, standards can also give rise to some competition law risks. These include facilitating unlawful information exchange, foreclosing innovative technologies and excluding certain companies from effective access to the established standard.¹⁹⁸

¹⁹⁶ *Id.*, ¶ 257.

¹⁹⁷ *Id.*, ¶¶ 263, 308, 325-332.

¹⁹⁸ Horizontal Cooperation Guidelines, *supra* note 94, ¶ 264 *et seq.*.

Researchers have suggested participating in standard setting.¹⁹⁹ We can view two types of proposed cooperation as standard setting: common performance measures and audit trails, on the one hand, and publication and release norms, on the other. One can structure both in ways that do not raise competition law concerns.

PAI is a particularly relevant body for AI companies. PAI collects best practices and sets informal standards. This process consists of PAI staff and “coalitions of the willing” amongst its members. This can involve setting up a steering committee to guide this process. It publishes outcomes on its website for public consultation, following which it considers responses, although how it takes responses into account is not entirely clear. Its executive committee, which is made up of profit and non-profit companies, oversees this process. There is no formal voting process at the steering committee or executive committee level. Its standards are non-binding, “best practice” principles, and compete with several other voluntary standards. However, given the size and importance of PAI’s members, it is likely that any standard—even if technically

¹⁹⁹ Peter Cihon, *Standards For AI Governance: International Standards To Enable Global Coordination In AI Research & Development*, FHI TECHNICAL REPORT 28 (Apr. 2019), https://www.fhi.ox.ac.uk/wp-content/uploads/Standards_-FHI-Technical-Report.pdf; *Ethically Aligned Design, First Edition: A Vision For Prioritizing Human Well-Being With Autonomous And Intelligent Systems*, IEEE REPORT (July 29, 2018), <https://standards.ieee.org/content/dam/ieee-standards/standards/web/documents/other/ead1e.pdf>.

non-binding—will have a significant market impact, that is, broad take-up in practice. Some risks may arise because there is some lack of transparency in how exactly the standard-setting process works. To address this, PAI could consider publishing a more formalised procedure on its website. Given the importance of the public consultation stage in ensuring fair participation in the standard-setting process, more clarity on how the process takes views into account is also advisable.

Other more formal SSOs in this space include IEEE, ICANN, the US National Institute of Standards and Technology (NIST), International Organization for Standardization (ISO), and the International Telecommunication Union (ITU). The IEEE, for example, has stringent standard-setting processes, which is published on its website.²⁰⁰ These include ensuring that all interested parties can participate; highly visible standard creation processes; formal voting procedure; and right to appeal a standard decision open to anyone at any point in time.

²⁰⁰ *Develop Standards*, IEEE SA, <https://standards.ieee.org/develop/develop-standards/govern.html> (last accessed: 25 August 2020).

*1. Suites of Tools or Guides*Key Recommendations

If the standard is voluntary and effectively competes with a number of other voluntary standards, it likely falls outside of Article 101(1) altogether and is the safest course. If not, participation in the standard-setting process, and the ability of third parties not involved in the standard-setting process to access and comply with the standard, should be on fair, reasonable and non-discriminatory (FRAND) terms.

Analysis

AI developers could create and share suites of tools and guides for machine learning that include measures of performance against collaboratively established standards.²⁰¹ This could be useful for fairness, safety, explainability, or robustness to adversarial inputs. In this section, we take standards around privacy-preserving machine learning as our example.

Standardized public benchmarks—typically a dataset and a performance target—are a key way to measure, guide and encourage progress in ML, from MNIST (handwriting)²⁰² and ImageNet

²⁰¹ Brundage, Avin, Wang *et al.*, *supra* note 4, at 25.

²⁰² Yann LeCun, Corinna Cartes, Christopher J.C. Burges, *The MNIST Database of Handwritten Digits*, <http://yann.lecun.com/exdb/mnist/> (last accessed: Aug. 25, 2020).

(image classification)²⁰³ to SQUAD (reading comprehension).²⁰⁴

Researchers can train their systems on these datasets and see how they perform: are they state-of-the-art? However, as benchmarks can guide progress, they can also affect in what domains progress is made and so the overall competitive landscape. For example, Amazon released logistics benchmarks²⁰⁵ and Audi released self-driving car benchmarks.²⁰⁶ There can be competing benchmarks. Chinese universities and companies created the AIBench²⁰⁷ benchmarking suite purportedly to balance MLPerf,²⁰⁸ which is largely developed by US universities and companies. PAI is developing benchmarks for safety, fairness, and other ethical objectives, for example through ABOUT ML²⁰⁹ and SafeLife.²¹⁰

²⁰³ Dawnbench, STANFORD DAWN, <https://dawn.cs.stanford.edu/benchmark/ImageNet/train.html> (last accessed Aug. 25, 2020).

²⁰⁴ Archa Jain, *Reading Comprehension With Squad*, STANFORD UNIVERSITY REPORTS, <https://web.stanford.edu/class/archive/cs/cs224n/cs224n.1174/reports/2758649.pdf> (last accessed Aug. 25, 2020).

²⁰⁵ Bharathan Balaji, Jordan Bell-Masterson, Enes Bilgin et al., *ORL: Reinforcement Learning Benchmarks for Online Stochastic Optimization Problems*, arXiv:1911.10641v2 (Dec. 1, 2019), <https://arxiv.org/pdf/1911.10641v2>.

²⁰⁶ Jakob Geyer, Yohannes Kassahun, Mentar Mahmudi et al., *A2D2: Audi Autonomous Driving Dataset*, arXiv:2004.06320v1 (Apr. 14, 2020) <https://arxiv.org/pdf/2004.06320>.

²⁰⁷ Wanling Gao, Fei Tang, Lei Tang et al., *AIBench: An Industry Standard Internet Service AI Benchmark Suite*, arXiv:1908.08998v2 (Aug. 12, 2019).

²⁰⁸ MLPerf, <https://mlperf.org/> (last accessed Sept. 7, 2020).

²⁰⁹ *ABOUT ML*, PARTNERSHIP ON AI, <https://www.partnershiponai.org/about-ml/> (last accessed Sept. 7, 2020).

²¹⁰ Carroll Wainwright and Peter Eckersley, *Introducing SafeLife: Safety Benchmarks for Reinforcement Learning*, NEWS IN PARTNERSHIP ON AI (Dec. 4, 2019), <https://www.partnershiponai.org/safelife/>.

One associated problem that standard tools and measures of performance face is that AI systems tend to lack a traceable log of steps in problem-definition, design, development, and operation. This could lead to a lack of accountability for subsequent claims about those systems' properties and impacts. Standard audit trails are a proposed solution for safety-critical applications of AI systems and refer to a traceable log of steps in system operation, and potentially also in design and testing.²¹¹

There are two relevant elements here: (1) collaboratively established standards, especially for privacy-preserving machine learning technologies and (2) sharing the tools and guides to comply with the standards. Both potentially offer clear pro-competitive benefits, including enhancing the privacy protection of machine learning technologies; helping companies to comply with standards; or helping to more consistently measure and audit the compliance of companies to these standards. However, one must weigh these benefits against any restrictive effects on competition on the markets that the standard may affect.

A standardization agreement should be analyzed in a three-step framework. First, it may not be capable of restricting competition and therefore fall outside of the Article 101(1) prohibition altogether if it meets certain conditions. If so, no further

²¹¹ Brundage, Avin, Wang *et al.*, *supra* note 4, at 24.

analysis is required. If it does not meet these conditions to bring the agreement outside Article 101(1), the agreement may be capable of restricting competition and therefore one should assess it under Article 101(1). Third and finally, even if the standardization agreement infringes Article 101(1) (if the previous conditions are not met) it can still be exempt under Article 101(3) if the standards are indispensable in producing efficiency gains that outweigh the restrictions on competition, those gains are passed on to the consumer, and the agreement does not eliminate competition on the market.

First, a standardization agreement will not generally be capable of restricting competition and hence fall outside the Article 101(1) prohibition if the standard does not risk creating market power²¹² or where a number of conditions are satisfied.²¹³ The analysis should be done according to the following potential relevant markets: (1) the market for the product to which the standard relates (for example, machine learning technology); (2) where the standard involves selection of separate technology (such as a specific type of software to ensure privacy standards), that related technology market; (3) the market for standard-setting itself, if different and competing standard-setting bodies or agreements

²¹² Horizontal Cooperation Guidelines, *supra* note 94, ¶¶ 277-278.

²¹³ *Id.*, ¶¶ 280-283.

exist; and (4) the market for testing and certification, where applicable.²¹⁴

Whether the standardization agreement could create market power—thereby increasing the competition law risk—will depend, amongst other factors, on the market shares of the technologies based on the standard.²¹⁵ It might not be possible to assess with any certainty at an early stage whether a large part of the industry will adopt a standard. In that case, one would use the market shares of the companies participating in developing the standard as a proxy. If PAI is leading any standard setting, though, it is likely that the standard-setting agreement will risk creating market power.

If the standard is liable to create market power, it may nevertheless fall outside Article 101(1) if the following conditions are satisfied: (1) members of the standard-setting organization or group are free to develop alternative standards or products that do not comply with the agreed standard; (2) all competitors on the market(s) affected by the standard²¹⁶ (see the four markets described above) can fairly participate in the process leading to the selection of the standard;²¹⁷ (3) the standard-setting process is open and transparent, with an objective and non-discriminatory procedure for

²¹⁴ *Id.*, ¶¶ 261-262.

²¹⁵ *Id.*, ¶ 296.

²¹⁶ *Id.*, ¶¶ 261-262.

²¹⁷ If fully open participation is not practicable, all players on the market should have sufficient representation in the standard-setting process and be informed of developments.

allocating voting rights;²¹⁸ and (4) access to the standard is on fair, reasonable and non-discriminatory terms (FRAND).²¹⁹

In the case of PAI, for example, any standard is likely to create market power, so these four conditions must be satisfied to fall outside Article 101. Its standards are voluntary and non-binding, and therefore its members are free to develop alternative standards or products that do not comply with the agreed standard. Access to the standard is also on FRAND terms as they are publicly accessible on its website. The process includes an unrestricted public consultation phase. This means all affected companies can fairly participate in the process. However, its standard-setting process is currently difficult to ascertain. To mitigate the competition law risk, the process could be more transparent. This is especially important given it does not have a formal voting or balloting system. To enhance transparency, PAI might consider publishing a clear standard-setting process guide on the PAI website, for example.

If a standardization agreement does not clearly comply with the conditions set out above, it should be assessed under Article 101(1), again along the (up to four) relevant markets discussed above. The aim is to ascertain whether the standard can limit differentiation and technical development. Assuming that companies did not design the standards to implement an anti-

²¹⁸ Horizontal Cooperation Guidelines, *supra* note 94, ¶ 281.

²¹⁹ *Id.*, ¶ 297.

competitive cartel, which would likely render it a “by object” infringement, the next step is to analyze the agreement under an effects-based analysis.²²⁰

Under an effects-based analysis, the first question is whether the standard is binding. A standard will be less restrictive of competition if the members of the standard-setting organization (SSO) remain free to develop alternative standards and products that do not comply with the agreed standard. In addition, standards that only cover minor aspects of the end product are less likely to lead to competition concerns compared to standards that are more comprehensive.²²¹ In the present case, the answer will likely depend on the intended use of the machine learning technology, and the importance of the privacy-preserving element to that use. However, PAI’s standards are non-binding and therefore will be less likely to restrict competition.

The greater the likely market impact of the standard, the more important it is to ensure equal access to the standard-setting process in order to reduce competition law risk. A standard established by an organization such as PAI with industry influence and participants covering a wide swathe of the market is likely to

²²⁰ A standard agreement that is used to implement a cartel, i.e., as part of a broader restrictive agreement aimed at excluding competitors will be a ‘by object’ infringement of Article 101(1). This does not apply to the coordination in question, which seeks to enhance consumer benefit by helping to develop privacy-preserving machine learning. Therefore, this analysis focuses on an effects-based analysis under Article 101(1).

²²¹ Horizontal Cooperation Guidelines, *supra* note 94, ¶ 293.

have a large market impact. At the same time, it may be impossible or extremely difficult for participation to be fully open to all competitors and/or stakeholders on all the relevant markets. In this scenario, it may be permissible to restrict participation by ensuring that the SSO keeps all stakeholders at least informed and consulted on the work in progress, even if they do not have a right to vote on the result.²²² PAI's public consultation process is helpful in this regard in ensuring wide participation, but it could be clearer on how these responses are given due regard. PAI also keeps stakeholders informed on the work in progress through regular blogposts, emails, and meetings, which helps to reduce the competition risk.

A standard-setting agreement should not clearly discriminate against any of the participating or potential members, either in the standard-setting process or the substance of the standard itself. For example, any standard-setting process that explicitly excludes new entrants, or companies active in certain countries, would be acting in a clearly discriminatory manner. In addition, both upstream and downstream companies that the standard affects should have an opportunity to participate in the standard. A blanket ban on either would be difficult to objectively justify. If manufacturers of soft drinks sought to establish a new standard on the packaging materials for use in their drink cans, for example, they

²²² *Id.*, ¶¶ 295-296.

should also include packaging suppliers in the upstream market, as the standard will impact them.²²³

Another example is if the technical specifications of the standard audit trail only allow use of a particular auditing software, even if it is of recognized equivalent performance to others. Such a standard would not satisfy the non-discriminatory condition and may reduce or prevent innovation because companies who have invested in an audit software may face large switching costs. There is also no objective justification or efficiency gains given the equivalence of performance.²²⁴

Finally, all competing undertakings should have fair access to the standard, including any intellectual property rights that they need to implement the standard. Transparency is a key element of ensuring fair access. If the requirements of the standard are unclear and imprecise, or simply not published to those who were not involved in the standard setting, this will have a foreclosure effect on third parties who will find it difficult or impossible to comply.²²⁵ However, where there are several competing standards or there is effective competition between standardized and non-standardized solutions, a limitation of access may not raise any competition

²²³ “Upstream” markets in a supply chain generally refer to material inputs needed for production, as opposed to “downstream” markets where goods get produced and distributed.

²²⁴ Horizontal Cooperation Guidelines, *supra* note 94, ¶ 325.

²²⁵ *Id.*, ¶ 324, by analogy.

risks.²²⁶ PAI “best practice” standards are fairly accessible to the public because PAI publishes them on its website.

Exemption under Article 101(3)

There may be some justifiable scenarios where it is not possible for the standards to comply with the conditions above: to be entirely open, accessible, and fully participative. Where this is the case, the standard may be liable to restrict competition in breach of Article 101(1), especially if one expects it to have a significant market impact. Nevertheless, the standard may satisfy the exemption under Article 101(3) if it generates countervailing efficiencies that outweigh the restrictive effects. The following analysis looks at each of the conditions for exemption in turn.

First, under the efficiencies condition, standards that enhance the comparability of privacy-preserving machine learning can produce clear efficiency gains. They can enhance competition between technologies on this front and therefore drive innovation by allowing companies to compete along agreed parameters. One could satisfy this condition by arguing that a privacy-preserving machine learning standard is beneficial to consumers because it improves the quality of the technology.

Second, to satisfy the indispensability condition, any restrictions of competition arising from the standardization

²²⁶ *Id.*, ¶ 294.

agreement should not go beyond what is necessary to achieve efficiencies. Where participation in standard setting must be limited to avoid serious inefficiencies in the process, the SSO should put recognized procedures in place to ensure there is a collective representation of interests.²²⁷ However, competition law would generally see a requirement for technical competency as a reasonable eligibility criterion for participation in standard setting.

The third requirement of Article 101(3) is that the efficiencies are passed on to consumers. In our case, it is likely the standardized measures and audit trails will facilitate competition between different technologies. It is therefore likely that efficiencies will pass-on to consumers.²²⁸

Finally, the standard should not eliminate competition in a substantial part of the market. If a standard is expected to become a *de facto* industry standard covering a wide part of the market, it must not foreclose third parties from effective access to the standard, thereby eliminating competition. This final condition relates to transparency—third parties must be able to clearly ascertain what is required in order to comply.²²⁹

²²⁷ *Id.*, ¶ 316. For example, in the Commission decision *X/Open Group*, the Commission held that if every company that was willing to commit itself to the Group objectives had a right to become a member, it would create practical and logistical difficulties for the management of work and possibly prevent appropriate proposals being passed (Case IV/31.458, *X/Open Group*, Commission decision of 6 February 1987, 1987 O.J. (L 35) 36, ¶ 45.

²²⁸ Horizontal Cooperation Guidelines, , *supra* note 94, ¶¶ 321-323.

²²⁹ *Id.*, ¶ 324.

Conclusion & Mitigating Steps

Generally, the proposal for standardized measurements for privacy-preserving ML and standard audit trails may fall outside Article 101(1) if one can show they provide scope for significant pro-competitive benefits. This is so long as the setting of the standard is on a fair and open basis, that the standard is voluntary to the extent possible, and that access to the standard is also on a FRAND terms.

Participation in standard setting can be subject to objectively justifiable criteria to ensure quality and efficiency (such as technical competence), but there should not be any unnecessary barriers to membership to the standard-setting process. To ensure that setting such a standard through PAI does not raise any competition concerns, PAI could be more transparent about its standard-setting process: especially on how the process gives public consultation responses due weight, and the decision-making process given the lack of a formal voting process. Finally, any exchange of competitively sensitive information during the standard setting process will also be subject to the usual rules on information exchange.²³⁰

²³⁰ See *supra* Section IV.B (Information Sharing).

2. *Publication and Release of Norms*Key Recommendations

The principles that apply to standardization on audit trails and comparability measures apply in a similar way to publication norms. Generally, norms that provide voluntary guidance or expert advice to help ensure safe and responsible publication will not breach competition law. However, more extreme norms that prescribe restricting or withholding publication altogether could give rise to competition law risks, and the potential for abuse.

Analysis

Given the potential harm from some AI systems falling into the hands of malicious actors, it is important to ensure that any “dual-use” (potentially dangerous) research that companies publish, and technology that they release, is made public in a responsible and considered way. This is the reason why PAI, for example, is currently coordinating discussions around publications norms to avoid such pitfalls.²³¹ Norms could be part of a self-assessment or voluntary regime, where firms must satisfy themselves that they comply before publication, and where rules act more as industry guidelines. Norms and guidance that PAI is currently considering, for example, include: (1) a comprehensive guide for researchers to

²³¹ *Publication Norms for Responsible AI: Ongoing Initiative*, PARTNERSHIP ON AI, <https://www.partnershiponai.org/case-study/publication-norms/> (last accessed Aug. 25, 2020).

self-assess their work; (2) access to a pool of experts who can advise on difficult publication decisions and facilitate consideration of risks; and (3) guidance on how to structure internal review processes designed to assess potential risks prior to publication.

A more radical—and therefore far less likely to be used—set of norms could recommend or require the withholding of some types of potentially dangerous research from the general public: limiting its publication or withholding publication altogether. This could be formalised through an industry contractual agreement to be bound by the publication norms. An alternative model would be an organization such as PAI needing to issue a formal approval or certification that a publication or release complies with the norms before a company makes it public.

As a preliminary point, any publication norms could only fall within Article 101(1) if they are agreed between competitors, for example through an association such as PAI. On the other hand, a purely unilateral policy of an undertaking, like an academic journal's publication policy or NeurIPS' impact assessment would likely not fall within the scope of Article 101(1).²³² The nature of the norms is also important in assessing the competition law risk. Publication norms that provide voluntary guidance or expert advice

²³² Note that a non-profit can also be an “undertaking” under Article 101(1), if certain conditions are satisfied. *See supra* Section II.A (Article 101(1) Prohibition).

should be analyzed in a similar way to privacy-preserving standards.²³³ They would usually be consistent with competition law as long as they are voluntary and effectively compete with a number of other voluntary standards.

However, risks could arise for norms that potentially prohibit the public release of certain research or new products altogether, even if one intends those norms to achieve a legitimate purpose. This could constitute an agreement between competitors to limit production or output, which may be an infringement “by object.” This means that it reveals, by its very nature, a sufficient degree of harm to competition that an analysis of its effects is not necessary to conclude that it breaches Article 101(1).²³⁴ Although it is theoretically possible for a “by object” restriction to benefit from Article 101(3), in practice any exemption argument will be difficult to bring successfully.

Furthermore, norms that seek to restrict publication of research could be more liable to abuse. For example, certain companies could skew the norms in a direction that would have a disproportionately restrictive effect on the publication or releases of a rival AI company or of new entrants or may go beyond what is necessary to achieve the legitimate purpose of the standards. This

²³³ See *supra* Section IV.C (Standard-setting).

²³⁴ Case C-67/13 P, *Groupement des Cartes Bancaires v Commission*, *supra* note 51, ¶¶ 49, 57.

might be a particular risk where certain AI companies have undue influence over the standard-setting process. It might also be a particular risk because of the potentially fluid and indeterminate line between what is a “potentially dangerous” publication and not. An agreement that was abused in this way would likely be treated by a competition authority as forming part of a broader restrictive agreement aimed at excluding competitors and may constitute a “by object” infringement of Article 101(1).²³⁵ It would also be difficult to justify under Article 101(3).

A deliberate abuse of the standard to restrict competition, for example to hinder new entrants, is unlikely to benefit from the exemption under Article 101(3). Outside of an abuse scenario, publication norms should be assessed under the Article 101(3) exemption in a similar way to those outlined above²³⁶. However, some types of publication norms, especially those that seek to restrict publication of research may be more challenging to satisfy compared to privacy standards for a number of reasons.

First, the EC generally sees standards that enhance “quality” factors such as quality, safety, and environmental aspects as beneficial to the consumer, and therefore giving rise to efficiency gains.²³⁷ Publication norms to, for example, prevent research being

²³⁵ Horizontal Cooperation Guidelines, *supra* note 94, ¶¶ 273-75.

²³⁶ *See supra* Section IV.C (Standard-setting).

²³⁷ Horizontal Cooperation Guidelines, *supra* note 94, ¶¶ 308-13.

used by malicious actors could constitute an increase in safety or security, and should therefore satisfy this condition in principle. Second, restrictions on release or publication must be indispensable to the achievement of those efficiencies and go no further. For extreme norms that prohibit certain product launches or research publications altogether, one should consider whether there is a less restrictive alternative that could still achieve the legitimate objective.²³⁸ These alternatives might include, for example, publishing to some authorized circles rather than the public as a whole. This can still allow consumers to benefit from the pro-competitive benefits of publication of research. Another, less restrictive alternative could be publishing in a partly redacted form, withholding the parts of the publication that are most liable to abuse, but releasing the remainder.

Third, the party seeking to rely on the exemption must demonstrate that countervailing efficiency gains that are passed on to consumers outweigh the restrictions on competition.²³⁹ The challenge here may be to demonstrate that the safety gains from keeping this research or technology away from malicious actors sufficiently outweighs the restriction on competition. It may be more difficult to argue pass-on for these types of efficiencies compared to, for example, standards on the battery life of fire alarms or

²³⁸ *Id.*, ¶¶ 314-19.

²³⁹ *Id.*, ¶ 321.

emissions from cars. Arguably, the benefits of the latter cases are more direct, tangible, testable and measurable.²⁴⁰ Fourth, if the norm is expected to become a *de facto* industry standard covering a wide part of the market, it must not foreclose third parties from effective access to the standard, thereby eliminating competition. In other words, third parties must be able to clearly ascertain what is required in order to comply.²⁴¹

Conclusion & Mitigating Steps

The principles that apply to standardization on audit trails and comparability measures apply in a similar way to publication norms. Norms that provide voluntary guidance or expert advice to help ensure safe and responsible publication will not usually breach competition law. However, more prohibitive norms that prescribe restricting or withholding publication altogether could give rise to competition law risks. These more extreme norms could be a “by object” infringement of Article 101(1) to the extent that they restrict publications or releases altogether. Furthermore, it could be challenging to rely on the Article 101(3) exemption. First, it may be difficult to demonstrate that there are consumer efficiencies that outweigh the restrictive effects given the difficulty in quantifying the “benefit” from the norms to the required evidential level.

²⁴⁰ This is similar to the challenge with justifying the Windfall Clause or Assist Clause, discussed *supra* in Sections III.A (OpenAI Assist Clause) and III.B (The Windfall Clause).

²⁴¹ Horizontal Cooperation Guidelines, *supra* note 94, ¶ 324.

Second, the “indispensability” element of the Article 101(3) exemption is particularly important here. The salient question is whether there is a less restrictive alternative that can still achieve the legitimate objective, such as publishing in redacted form or to a trusted circle, rather than withholding publication altogether. To mitigate competition law risks, the least restrictive alternative should always be the preferred avenue in the first instance.

Another risk from these extreme norms is the possibility that companies could use them to deliberately dampen competition. An important safeguard is to ensure that a fair and open system of participation exists and that no one company has undue influence throughout the standard-setting process. For these reasons, more voluntary forms of guidance may be advisable as lower risk alternatives to these more extreme norms.

CONCLUSION

In this paper, we discussed how the competition law prohibition on restrictive agreements under Article 101(1) might apply to forms of cooperation between AI companies, advised on the potential level of risk and suggested practical steps to mitigate any risks. AI cooperation seeks to ensure that AI is developed and deployed safely, securely, and beneficially. Forms of AI cooperation range from AI bias auditing and setting safety benchmarks, to those that seek to address longer-term risks of catastrophic accidents or

power concentration that may arise if and when we near the development of TAI, like the Assist Clause and Windfall Clause.

Whilst the goals of AI cooperation are crucial and legitimate, cooperation between competing AI companies can risk infringing Article 101(1), for example, where they involve an agreement or understanding to restrict competition between AI companies or the exchange of commercially sensitive information. The scope of Article 101(1) is broad, ranging from binding agreements to scenarios where an undertaking “knowingly substitute practical cooperation between them for the risks of competition.”²⁴² An agreement can breach Article 101(1) both where it is capable of restricting competition and where it actually does so. The prohibition may apply as soon as undertakings enter into an agreement, without any need for them to (ever) implement the agreement.

However, this Article argued that some relatively simple and “low-cost” mitigating steps that do not sacrifice the objectives of the AI cooperation strategies could significantly help to reduce competition law risks. These mitigating steps may restructure the agreement so that it is not capable of producing anti-competitive effects or by ensuring that an AI company acts unilaterally rather than cooperatively whilst retaining the gains from an AI governance

²⁴² Case C-48/69, *Imperial Chemical Industries v European Commission (Dyestuffs)*, 1972 E.C.R. 619.

perspective. These may help to structure cooperation outside of the Article 101(1) prohibition altogether.

Even if one cannot easily restructure the cooperation to fall outside Article 101(1), mitigating steps might allow the cooperation to benefit from the exemption under Article 101(3), if one can demonstrate countervailing efficiencies and the agreement is necessary to achieve them. In relation to responsible AI development, the AI company seeking to rely on the exemption would need to demonstrate efficiency gains from the development of an improved AI system, for example, one that has been more rigorously tested for ethics or safety, to the benefit of consumers.

There will be scenarios where one cannot so easily restructure cooperation strategies to comply with competition law. This presents a more difficult trade-off between gains in responsible AI development and the risk of breaching competition law. One example may be implementation of the Assist Clause at a relatively late stage in TAI development where its implementation may lead to a significant reduction in competition. In these instances, we would recommend seeking specialist competition law advice as the level of risk in each case would depend on several specific factors. These include the market position of the parties, how one defines the relevant market, the level of competition in the rest of the

market, and the strength of evidence of any countervailing efficiencies for the purposes of an Article 101(3) defense.

Overall, EU competition law should be an important consideration for all AI companies and AI governance researchers when planning AI cooperation strategies. Substantively, AI governance and competition law share similar goals: to prevent concentrations of power, to encourage innovation, and to promote societal and consumer welfare. Prudentially, the EC's current regulatory focus on the tech sector, as well as its strategic focus on AI and Big Tech's activities in AI development, are likely to continue for the foreseeable future. This makes it more likely that the EC will bring competition enforcement against AI companies. Sanctions can be wide-ranging and significant, as well as expending significant company resources in defending investigations that can go on for many years.

The stakes are high on both sides. Competition law is important and powerful; AI cooperation seeks to avert potentially catastrophic harms. However, EU competition law should not be an insurmountable obstacle to AI cooperation strategies that legitimately seek to enhance safe and responsible AI development and deployment. This paper advocates for AI companies to build a good awareness of the key principles applicable to AI cooperation, to maintain ongoing diligence to ensure they are structuring forms

of AI cooperation in ways that are consistent with competition law, and to seek specialist competition law advice where relevant. These steps can significantly help reconcile the tensions between cooperative AI strategies and EU competition law.

Appendix 1: Table of Implementation Scenarios for the Assist Clause

Implementation Method	Potential Risks
<p>Full Acquisition: Competitor is fully acquired by Leader</p>	<p>This acquisition will likely be subject to merger review, assuming turnover thresholds are met.</p>
<p>Partial acquisition: A part of the competitor that is developing TAI-relevant technology (e.g. safety, language models, reinforcement learning environments) is acquired by the Leader, the rest of the competitor continues to operate on the market</p>	<p>The principles in (1) above also apply here.</p> <p>However, may be lower risk than full acquisition to the extent there are fewer ‘overlaps’ in a partial acquisition i.e. relevant markets where the competitor and Leader overlap.</p> <p>This should reduce the level of antitrust scrutiny and therefore reduce the likelihood that the merger will be blocked/ subject to remedies.</p>
<p>Joint venture (“JV”) with Leader: parties enter into a joint commercial enterprise, but otherwise parties retain their distinct identities</p>	<p>The competition law analysis would depend on whether the JV would be a ‘full-function joint venture’ (FFJV). A FFJV would be notifiable under EUMR, whereas a non-FFJV would not be notifiable but would be assessed under Article 101. It is preferable for a JV not to be notifiable under the EU merger regime, because it will be subject to a higher level of scrutiny compared to under Article 101(1).</p> <p>As discussed above, it is advisable to structure the JV to avoid being caught under EUMR. This could be done by ensuring that the JV only takes over one specific function of its parents' activities, such as R&D, which would mean the second condition for a FFJV is not satisfied. Or it could be made clear that the JV is only intended to be established for a short, finite period and not meant to be long-lasting (fourth condition).</p> <p>However, a non-full function JV would still need to be assessed under Article 101(1).</p>
<p>Partial joint venture with Leader: i.e. only contributing the part of the competitor that is developing TAI-relevant technology</p>	<p>The legal principles set out in row (3) above also apply here.</p> <p>In practice, a partial JV may be less likely to qualify as a FFJV that is notifiable under EUMR if e.g. the JV is only taking over one specific function of the parent, or does not have its own dedicated resources. However, this will depend on what resources/functions the Leader contributes to the JV, as the conditions must be satisfied by considering the JV as a whole.</p> <p>Even if a partial JV is notifiable, there are potentially fewer ‘overlaps’ compared to a ‘full’ JV i.e. relevant markets where the competitor and Leader overlap. This will reduce the level of antitrust scrutiny and therefore reduce the likelihood that the merger will be blocked.</p> <p>If a partial JV is not notifiable, it would still need to be assessed under Article 101(1).</p>
<p>“Unilateral” R&D slow-down: competitor unilaterally slows down its R&D on TAI</p>	<p>If the competitor can establish that it is slowing down its R&D on a genuinely unilateral basis, then it is likely to fall outside of the Article 101(1) prohibition.</p> <p>However, the R&D slow-down could constitute a form of cooperation subject to Article 101(1) if it implements a collaborative Assist Clause. In this scenario, the</p>

	<p>slow-down could be seen as part of a single and continuous infringement flowing from the Assist Clause itself.</p> <p>In a scenario where the competitor acts under a unilateral Assist Clause, the slow-down could still breach Article 101(1) if it involves any agreement, understanding or exchange of information at the point of implementation (e.g. any coordination with the Leader as to the timing or method of the slow-down).</p>
<p>R&D shift: competitor shifts R&D to a particular topic (e.g. safety or applied/ product development)</p>	<p>If the shift in R&D is truly unilateral, then it would fall outside of the Article 101(1) prohibition.</p> <p>However, this R&D shift may be a breach of Article 101(1) if it implements a collaborative Assist Clause, wherein the R&D shift and the Assist Clause would form part of a 'single and continuous' infringement. The potential overall infringement could be seen as an agreement between parties that are part of the pattern of collaborative Assist Clauses to restrict competition in the market for TAI development by shifting TAI R&D development.</p> <p>However, it may be possible to rely on the Article 101(3) exception if the parties can demonstrate that the switch to safety R&D generates efficiencies that will benefit consumers, and those efficiencies outweigh the restriction on TAI development competition.</p> <p>Note also that the competition analysis is per product market, therefore the countervailing efficiencies must be in relation to TAI, and not any other product, in order to fulfil the conditions under Article 101(3).</p> <p>In a scenario where the competitor acts under a unilateral Assist Clause, the shift could still breach Article 101(1) if it involves any agreement, understanding or exchange of information at the point of implementation (e.g. any cooperation with the Leader as to the timing or nature of the shift).</p>
<p>Competitor stops TAI R&D, and shifts to applied AI development (e.g. a product development)</p>	<p>See (6).</p>
<p>Staffing changes: Competitor puts many employees on gardening leave or fires them.</p>	<p>If the competitor can establish that it is really acting unilaterally, then this is likely to fall outside the Article 101(1) prohibition.</p> <p>In practice, it may be difficult to show that putting many employees on gardening leave/ firing them, which effectively restricts the competitor's R&D efforts and potential output, is not part of implementing a collaborative Assist Clause, i.e. an agreement or understanding between competitors. If so, the Assist Clause and staffing changes would form part of a single and continuous infringement.</p> <p>The difference between this option and the R&D switch is that this is less likely to satisfy Article 101(3)- there is no direct countervailing benefit from alternative forms of R&D. We would need to rely on the more indirect benefits from enhanced safety development- a more challenging argument to bring for reasons discussed above.</p> <p>In a scenario where the competitor acts under a unilateral Assist Clause, the shift could still breach Article 101(1) if it involves any agreement, understanding or exchange of information at the point of implementation (e.g. any cooperation with the Leader as to the timing or nature of the staffing changes).</p>

<p>Employee resignations: competitor's employees resign or are fired</p>	<p>If the competitor's employees are fired, the same analysis for scenario (6) and (7) above apply i.e. the Assist Clause has the effect of restricting competition in TAI development by prompting/obliging the competitor to take this action. If it can be shown to be truly unilateral, then no issues arise.</p> <p>As an initial point, although only undertakings can be bound by competition law, an undertaking will have strict liability for any actions of their employees under competition law, even if the employer has no knowledge of these actions or if the employee is acting contrary to the employer's instructions.</p> <p>If the competitor's employees resign, and it can be shown that they took this decision independently and without any influence or pressure from the competitor company, this may be lawful because it is a unilateral conduct, applying the same principles as discussed above.</p> <p>However, it is possible that the employees were influenced by the Assist Clause, even if they are not strictly bound by it. Therefore, it is arguable that the resignations are an implementation of the Assist Clause. In this case, the Assist Clause could be seen as an agreement that restricts competition, albeit indirectly, (by influencing the employees to resign in order to curb competition in TAI development). The Clause, as well as its subsequent implementation via the employee resignations, could therefore be treated as part of a single and continuous infringement.</p>
<p>Competitor's employees resign or are fired, and are hired by Leader</p>	<p>The competition law risk here will depend on what basis the firing and hiring / resigning and hiring occur, and whether these acts are truly unilateral (which would fall outside Article 101(1), or whether they are in fact implementing the Assist Clause.</p> <p>If this conduct is subject to an agreement or understanding between the competitor and Leader that the competitor's employees will be fired and the Leader will hire them, that could be an implementation of the Assist Clause that breaches Article 101(1). Here, the Clause and its subsequent implementation constitutes a single and continuous infringement that indirectly restricts competition (between the competitor and Leader in TAI development as a result of transfer of employees).</p> <p>In a scenario where the competitor acts under a unilateral Assist Clause, the hiring/firing could still breach Article 101(1) if it involves any agreement, understanding or exchange of information at the point of implementation (e.g. any cooperation with the Leader with respect to the hiring/firing).</p>
<p>Some employees (e.g. safety, language models, or reinforcement learning environments teams) resign or are fired, and are hired by leader</p>	<p>Principles from (10) will apply similarly here. See row above.</p>
<p>Competitor provides compute cheaply or free to leader</p>	<p>In this scenario, the competitor is providing the Leader with an input that is the main cost of AI R&D. This conduct may give rise to competition risks where the provision to the Leader is not unilateral, and it means that the competitor can no longer compete with it, or needs to lower output, reduce its own innovation, etc. For example, if the competitor never reached critical mass to be able to viably compete with the Leader, and would have dropped out of the market anyway, then this grant of compute would not have led to any appreciable reduction in competition that the competitor would otherwise be able to exert.</p>

	<p>However, just conferring an advantage to a competitor is not a breach of competition law per se. It depends overall on the effects on competition as a result of the provision on both parties. If the agreement to provide the compute is capable/intended to/ does indeed hamper the ability of the competitor to compete with the Leader, that would give rise to competition law breach (per the reasoning in (5)). If the grant of compute can be justified e.g. because the competitor wishes to switch to another area of AI R&D that requires less compute, and still brings countervailing benefits in e.g. AI safety, that may be compliant under similar reasoning to the R&D 'shift' in (6).</p> <p>Further, if the grant allows the Leader to gain a critical scale in input that speeds up or improves the quality of its AI R&D, that could be seen as a countervailing benefit that justifies the conduct under Article 101(3) as long as it outweighs the restriction on competition that the competitor is able to exert and does not go beyond what is necessary.</p>
Competitor provides research/ software cheaply or free to leader	Similar principles to above.

Appendix 2: Cooperation Strategies & Antitrust Analysis

Strategy	Description	Antitrust Concerns	Mitigating Steps	Discussion
Assist Clause	A central risk from TAI development is race dynamics leading to corner-cutting on responsible development. To address this, OpenAI proposed the 'Assist Clause' in its Charter, which commits it to stop competing and start assisting any 'value-aligned' company that gets close to developing AGI.	Entering into the Assist Clause could be an agreement to restrict competition; implementing the Assist Clause could also separately restrict competition.	Companies should not seek to agree with, persuade or influence their competitors to also adopt Assist Clause; implement assistance under the Clause in ways that are efficiency-enhancing, do not reduce overall output or innovation.	Section III.A
Windfall Clause	An agreement entered into by an AI company to redistribute a part of its wealth if and when its earnings exceed a certain threshold.	'Agreement to agree' between AI companies to enter into Windfall Clause could raise risks to extent Windfall Clause disincentivizes competition between companies, restricts sales or output etc.	Avoid entering into an 'agreement to agree' to the Windfall Clause; structure Windfall Clause to minimise disincentive effect.	Section III.B
Agreements on secure enclave	A set of software and hardware features that together provide an isolated execution environment that enables a set of strong guarantees for applications running inside the enclave.	Competing AI companies that enter into an agreement on secure enclaves could constitute a horizontal agreement on standards that could exclude competitors. See (12) on standardized benchmarks.	See (12) on standardized benchmarks.	Section IV.A
Mutual monitoring	AI companies could monitor or audit the responsible AI development of other AI companies.	Risks arise if competitors exchange commercially sensitive information between themselves.	Use a third party to carry out audits instead of a competing AI company; minimise the commercially sensitive information exchange to only what is strictly necessary for the immediate purpose.	Section IV.B.1
Red-teaming	To enhance AI safety, 'white hat' hackers try to uncover vulnerabilities in systems and organisations and then tell the people that run them.	See (4) on mutual monitoring.	See (4) on mutual monitoring.	Section IV.B.2
Incident sharing	AI companies share information about what has gone wrong with safety testing.	See (4) on mutual monitoring.	See (4) on mutual monitoring.	Section IV.B.2
Compute accounting	An AI lab should comprehensively account for the computing power	See (4) on mutual monitoring.	See (4) on mutual monitoring.	Section IV.B.2

	used in a major AI project, and share lessons learned on the challenges of precision and comparability across projects.			
Communication and 'heads-up'	<p>Greater communication between competitors could include regular updates, clear communication channels, and joint events such as workshops, retreats, conferences, residencies, country visits.</p> <p>Separately, pre-publication or pre-release 'heads-up' to competitors about R&D or products are intended to give competitors an opportunity to point out legitimate reasons why the R&D product should be stalled e.g. because of outstanding safety concerns, or because it can be exploited by malicious actors.</p>	See (4) on mutual monitoring. Risks also arise to the extent the 'heads-up' gives competitors the ability to restrict or delay output.	See (4) on mutual monitoring. Moreover, do not give competitors who are receiving the 'heads-up' the ability to restrict or delay output.	Section IV.B.3
Seconding staff	Seconding engineers or researchers from one AI lab to another to help trust building.	See (4) on mutual monitoring.	See (4) on mutual monitoring.	Section IV.B.4
Third-party auditing	Use a third-party trade association or safety standards body to carry out independent, third party auditing of AI systems.	Do not use the third party as an intermediary for competing firms to exchange the same commercially sensitive information, just indirectly.	Third-party audit should be on an independent basis, without disseminating any commercially sensitive information between competitors.	Section IV.B.5
Bias and safety bounties	Organizations commit to giving individuals recognition or compensation for recognizing and reporting bugs, especially those related to exploits and vulnerabilities. Organisations could establish a similar system for bias and safety problems around AI systems.	See (4) on mutual monitoring. Risks if the bounty hunter discloses confidential information about its technology or R&D in reporting how its systems exposed the vulnerability.	See (4) on mutual monitoring.	Section IV.B.6
Standardized benchmarks	These could include measures of performance against collaboratively established standards. This could be useful for fairness, safety, explainability, or robustness to adversarial inputs.	Risks where the standard has the potential to exclude competitors from the standard-setting process or from access to the standard once set.	Standard-setting process and access to the final standard should be on fair, reasonable and non-discriminatory (FRAND) terms.	Section IV.C.1
Standardized audit trails	Standardized audit trails are a subset of standardized benchmarks. They constitute a traceable log of	See (12) on standardized benchmarks.	See (12) on standardized benchmarks.	Section IV.C.1

	steps in system operation, and potentially also in design and testing. This helps to ensure that there is accountability for subsequent claims about safety-critical AI system's properties and impacts.			
Publication and release norms	Given the potential harm from AI systems falling into the hands of malicious actors, it is important to ensure that any 'dual-use' (potentially dangerous) research that companies publish, and technology they release, is made public in a responsible and considered way.	See (12) on standardized benchmarks. Additional risks where more extreme norms prescribe the restriction or withholding publication altogether, which could constitute an anti-competitive agreement.	See (12) on standardized benchmarks. Opt for more voluntary forms of guidance or expert advice; and avoid more extreme norms that restrict publication altogether.	Section IV.C.2

Appendix 3: Forms of Mutual Monitoring

Strategy	Competition Law Risk	Mitigation
<p>(1) Shared red teaming</p> <p>AI developers should conduct “red-teaming” (adversarial attacks) exercises. They are “white hat” hackers who try to uncover vulnerabilities in systems and organisations and then tell the people that run them.</p>	<p>The level of risk will depend on whether the red teams exchange any commercially sensitive information. The lowest risk approach is for the attacks to avoid sharing any commercially sensitive or non-public information.</p> <p>However, it may be necessary to share certain commercially sensitive information to achieve the safety-enhancing purpose of the red teaming. For example, the ‘hackers’ may wish to share technical information about their AI systems with the company that is subject to the adversarial attacks. This may be necessary to inform the subject company of the type of technology that was capable of exposing the vulnerabilities.</p> <p>Risks may be manageable if mitigating steps are taken.</p>	<p>Minimize information exchanged to what is necessary to achieve the purpose.</p> <p>Add safeguards to ensure staff only disclose the information within a ‘clean team’ if possible and only use it for the designated purpose.</p> <p>One can further mitigate risks if a third party and not competing AI labs carry out red teaming (unless there are clear justifications, e.g. third party has the technological acumen to carry out successful ‘hacks’).</p> <p>However, one can reduce risks to a manageable level just based on steps 1 and 2 above.</p>
<p>(2) Incident sharing</p> <p>AI developers should participate in “incident sharing” more often, including through multilateral channels (e.g. PAI).</p>	<p>If ‘incident sharing’ is publicly available, this should not give rise to any competition law risk.</p> <p>If incident sharing will only be carried out e.g. within the confines of a trade association such as PAI, it will be relevant to consider whether it would involve sharing any strategically useful information with other companies.</p> <p>If incident sharing only involves sharing past (historic) information about what has gone wrong with safety testing, that information would only be strategically significant if it nonetheless provides an insight into the future development of the company’s technology or R&D, e.g. by providing insight into its software or algorithms. If it does not provide such an insight, its disclosure is unlikely to raise competition law risk.</p> <p>However, if the incident sharing does provide a strategic insight, Article 101(1) potentially prohibits it. For example, it might be strategically significant in nature because it gives insight into R&D and essential technology unless information is historic or publicly available. However, incident sharing is probably only worthwhile if the incidents are sufficiently recent that they can be useful for parties to feed into their current safety development.</p> <p>To ensure one can rely on the Article 101(3) exemption, the commercially significant information should be the minimum necessary to achieve the legitimate purpose of enhancing safety testing. However, it is arguably an efficiency-enhancing measure (condition 1 of Article 101(3)) that allows competitors to benchmark against safety standards in the market.</p>	<p>Incident sharing should be released to the public domain if possible.</p> <p>If non-public, share only what is strictly necessary to achieve the legitimate objective.</p> <p>Share in an aggregated and/or anonymised form if possible through a third party such as PAI.</p>

	<p>A multilateral channel such as PAI could also collate the incident sharing. However, using PAI as a sharing mechanism will only reduce the competition law risk if they collect the individual incident data and share it with PAI members on an aggregated and anonymised form. Further, if having access to this incident sharing confers a material competitive advantage to members of the PAI, its membership criterion should be sufficiently open to all industry participants in order to comply with Article 101.</p>	
<p>(3) Compute accounting</p> <p>An AI lab should comprehensively account for the computing power used in a major AI project, and share lessons learned on the challenges of precision and comparability across projects.</p>	<p>Sharing compute data is likely to be commercially sensitive information because it is an essential input cost of an AI lab. If this is historic information, this is likely to be commercially sensitive only to the extent that even such historic information on compute power could provide an insight into the AI lab's future strategic direction.</p> <p>If this historic compute information does provide such insight, then exchanging such information between competing AI companies could breach Article 101(1). As above, consider instead using less anticompetitive options to achieve the legitimate objective of enhancing safety. For example, consider using a third party e.g. PAI to aggregate and anonymous computer information and then share that 'sanitized' information with the AI companies. If having access to this compute power information confers a material competitive advantage to members of the PAI, its membership criterion should be sufficiently open to all industry participants in order to comply with Article 101.</p> <p>If it is not possible to achieve the legitimate safety objective by this method e.g. because information must be individualised to be meaningful, then the question in each case should be 'is this information running the minimum competition law risk possible to still achieve the legitimate objective'?</p> <p>As to sharing the lessons learned about precision and comparability across projects, this could arguably be objective/technical information that has a pro-competitive purpose in allowing companies to benchmark against each other. It seems unlikely to have the object or effect of restricting competition.</p>	<p>See (2).</p>