

Artificial Intelligence and Trade

*Anupam Chander**

A. INTRODUCTION

When Lily Leong stepped outside that morning on July 1, 2025, the voice in her ear guided her to the nearest Lime ebike, only two blocks away. Her work was fifteen kilometers away in Jakarta's business district and her Samsung Universe One had woken her that morning, timing its gentle intrusion based on her sleep cycle. Her phone had reported that it was a good day to bike to work and had run through the day's appointments. As she walked, her Bose headset would gently interrupt her latest K-Pop favorite, Girls Next Generation, to tell her which way to turn ('Right after the Starbucks'). She was hoping to be able to save enough money by the end of the year to buy the Bose AR Glasses that would show her route without interrupting GNG's 'In a Funk'. On her ebike, the voice guided her around the construction site building a new skyscraper. She saw the Komatsu robot erecting the steel girders that framed the building. The construction site was marked as a Human Exclusion Zone, an 'HEZ,' with prominent signs depicting a diagonal line crossing out a human being. Humans supervised from a protected shelter across the street, staring at screens that connected them to cameras and robots. She stopped the ebike to frame a photo with an idle human in the foreground and the robot construction worker lifting a heavy steel beam in the background and uploaded it to Instagram.

As she arrived at the skyscraper where she worked, the glass turnstile whisked open, a screen displaying the photo from her first day at work two years earlier when she had long hair. In the elevator she put her hand to her mouth to muffle her laugh at the latest fad on TikTok – the #PetTwin challenge, where people showed their pets wearing hairstyles and clothes matching themselves using images generated by an app. Coming to her standing desk somewhere among the hundreds of desks on the

* Anupam Chander is Professor of Law at Georgetown Law. Contact: ac1931@georgetown.edu. The author thanks Sandeep Chandy for excellent research assistance.

fortieth floor, she sent a text to Xiaoice in Chinese about some issue she was having with her loud neighbor at work, and the Microsoft AI responded with suggestions on how to politely indicate her concern. Her Lenovo computer identified her through an iris scan, and a program automatically queued up her first task for the day – an appeal of the bank’s automated denial of a housing loan in Germany.

Invisible strings pulled by invisible computers across the world shaped Leong’s morning. Her Samsung phone relied on computers in Seoul to awaken her with useful information about the day. The voice telling her which turn to make for a safer biking route was Google’s Singapore computer. A Bose computer in Massachusetts played songs that it thought she would like. The Komatsu heavy machinery installing the steel girders and pouring the concrete was guided by Nvidia AI based out of Santa Clara, California, coordinating with Komatsu computers in Tokyo. Instagram’s California computers promoted her photo to followers, after scanning it for illegality. The facial recognition system was the work of Hikvision operating through computers in Shenzhen, China. The TikTok videos on her phone were selected for her by the Shanghai-based enterprise using leased Amazon servers in the United States. Microsoft ran its Xiaoice chatbot out of Beijing. The AI making the initial credit decision lived on Ping An Technologies’ servers in Shenzhen. Even less visible were the various smart city sensors and actuators operated by various unnamed companies in China, the United States, and Singapore – these systems operated the traffic signals, routed the garbage trucks, and deployed city resources.

Even if this scenario imagines the near future, the technologies mentioned largely exist today. Artificial intelligence (AI) is already crossing borders, learning, making decisions, and operating cyber-physical systems.¹ It underlies many of the services that are offered today – from customer service chatbots to customer relations software to business processes. AI is already powering trade today.

This chapter considers AI regulation from the perspective of international trade law. Because of the near-universal reach of trade rules, the focus here will be on the World Trade Organization (WTO) agreements. My argument unfolds as follows. Section B argues that foreign AI should be regulated by governments – indeed that

¹ The WTO describes AI as follows: ‘One way to look at AI is as the latest form of automation. However, instead of substituting machine power for manual labour, as in the past, the use of AI involves substituting the computing ability of machines for human intelligence and expertise. Human abilities that were once thought to be out of the reach of machines, such as making a medical diagnosis, playing chess or navigating an automobile, are now either routine or well within reach. Two uses of AI – analogous to the weak AI and strong AI distinction – may be distinguished here, i.e. AI which aids the production of goods and services, and AI which helps to generate new ideas. Examples of the former use of AI include guiding robots in warehouses, optimizing packing and delivery, and detecting whether loan applicants are being truthful. Examples of the latter use of AI are analysing data, solving mathematical problems, sequencing the human genome, and exploring chemical reactions and materials.’

See WTO, *World Trade Report 2018: The Future of World Trade: How Digital Technologies are Transforming Global Commerce* (Geneva: WTO, 2018), at 30 (references omitted).

AI must be what I will call ‘locally responsible’. Section C then refutes arguments that trade law should not apply to AI at all and shows how the WTO agreements might apply to AI, using two hypothetical cases – a medical diagnostic AI-based system and an insurance coverage decision-making AI. The analysis will reveal how the WTO agreements leave room for governments to insist on locally responsible AI, while at the same time promoting international trade powered by AI.

B AI’S KANGAROO PROBLEM, OR WHY REGULATE AI?

In 2018, President Emmanuel Macron announced that France will send regulators to sit inside Facebook to evaluate how the company combats hate speech on its services.² The regulators will meet with Facebook decision-makers not only in its offices in France, but in Facebook’s offices in Dublin, Ireland, and Menlo Park, California.³ President Macron called this ‘smart regulation’ and hoped to extend the model to the rest of ‘GAFA’ members – Google, Apple, and Amazon.⁴

But what about decisions made by AI? Indeed, while it has hired legions of human content moderators, Facebook is also depending on AI to make content moderation decisions. When Mark Zuckerberg testified before Congress in 2017, he cited ‘artificial intelligence’ more than thirty times in his deposition.⁵ ‘Over the long term,’ Zuckerberg offered, ‘building AI tools is going to be the scalable way to identify and root out most of this harmful content.’⁶ So, just as it may be appropriate for France to demand that Facebook’s human decision-makers in Ireland or California comply with its laws – at least with respect to information destined for France – it is appropriate for France to demand that Facebook’s AI decision-makers follow its laws on hate speech.

Governments have good reasons to regulate trade powered by AI. Imagine a dystopian turn to the sci-fi scenario in the introduction: your phone is listening in without permission and pushing advertising based on what it hears, your music app is selling your movements, the robot builder builds an insecure structure, the social network’s algorithms promote hate speech because they engender more engagement, the chatbot starts giving dangerous medical advice, the credit decisions are racially discriminatory, or the smart city is a massive surveillance system in the service of a repressive government.

² M. Rosemain, M. Rose, and G. Barzic, ‘France to ‘Embed’ Regulators at Facebook to Combat Hate Speech,’ *Reuters*, 12 November 2018, available at www.reuters.com/article/us-france-facebook-macron/france-to-embed-regulators-at-facebook-to-combat-hate-speech-idUSKCN1NH1UK.

³ *Ibid.*

⁴ *Ibid.*

⁵ D. Harwell, ‘AI Will Solve Facebook’s Most Vexing Problems, Mark Zuckerberg Says. Just Don’t Ask When or How,’ *The Washington Post*, 11 April 2018 (noting Zuckerberg’s promotion of artificial intelligence in connection with decisions related to removing speech for fake news, hate speech, discriminatory ads, and terrorist propaganda).

⁶ *Ibid.*

With respect to the broad array of services now increasingly powered by AI, there are many legitimate (by which I mean non-protectionist) reasons why a government might seek to regulate the underlying AI. AI operates quite differently from human beings, raising both new issues and also old issues in a new way. AI operates at a different scale, using a different evaluation process, without emotion and judgment. Some may see being subject to decisions taken by AI as an attack on their dignity, while others may worry about who will be held accountable for AI decisions.⁷ Regulations built for a world of human reasoning, emotion, and judgment may not equal a world where decisions are made by AI.

How is automated decision-making different? First, and obviously, it is done by computers rather than humans, and thus lacks traditional qualities of human judgment, empathy, and emotion, though it might offer facsimiles of any of these qualities. Second, the ability to transmit real-time data has enabled far more personalized cross-border decision-making than ever before – whether by humans or AI. Third, because it is computerized, it may be done at enormous scale. Fourth, while AI might not be programmed with invidious bias, it might learn that bias from the real-world data it receives – without even knowing perhaps to be mindful of the possibility of such bias.⁸

Decision-making from abroad, of course, predates the rise of AI. Banks, credit card companies, insurance companies, and the like have long relied on decisions made abroad. While there is nothing per se novel about decision-making or information processing across borders, the fact that the Internet now touches almost all of our daily activities increases the opportunities for AI-based decision-making, including decision-making across borders. AI changes the nature, scope, and scale of foreign decision-making. We are entering into a world in which your credit, your job prospects, your insurance claim, the news you read, and even the dates you go on are determined by faceless computers in a distant land.

There is a reason to believe that AI systems will make more mistakes as they cross borders. First, AI might be designed for different environments, nurtured on data from polities that might behave differently. This is a form of the well-known problem that AI trained on, say, a largely white (and male) population, might perform poorly with respect to other populations. Imagine, for example, an AI trained to recognize threats in the United States, but which fails to understand the context of threats in Myanmar – to possibly tragic consequences. Second, because of immense commercial pressures to claim the first mover advantage – attracting both media and venture capital, AI is being rolled out before it is ready. Because machine learning systems benefit from larger datasets, the opportunity to engage more people across the globe will tempt companies to apply their systems ever more broadly. Third, the quality of

⁷ M. E. Kaminski, 'Binary Governance: Lessons from the GDPR's Approach to Algorithmic Accountability,' *Southern California Law Review* 92 (2019), 1529–1616 (identifying concerns animating calls for regulating algorithmic decision-making).

⁸ *Liu v. Uber Technologies Inc.*, 20-cv-07499, District Court, N.D. California; A. Chander, 'The Racist Algorithm?', *Michigan Law Review* 115 (2017), 1023–1045.

AI's judgments will be hard to assess because firms have incentives to proclaim the effectiveness of their AI while individual users cannot amass the overall data necessary to evaluate it. Like the problem of legal transplants – which can prove unsuited in new social, cultural and legal contexts – AI transplants might prove problematic.⁹

Thus, there may be special reasons to distrust foreign AI, which may not have been trained on local conditions. I call this 'AI's Kangaroo Problem' in reference to the Volvo case, where Volvo realized that its 'Large Animal Detection' system initially failed to recognize kangaroos because of their jumping, and then began training its system with films of 'kangaroos' roadside behaviour.¹⁰ When a Tesla, apparently on autopilot, slammed into a stopped tow truck on a Russian road, one news account offered a conjecture: "Tesla cars [may not be] trained on Russian roads and vehicles."¹¹ More generally, AI will often need to be culturally or environmentally sensitive and an AI 'trained' on the behavior of the US population may well produce erroneous results when applied in China, or vice versa.

AI's Kangaroo Problem makes it especially urgent for governments to monitor foreign AI. Of course, higher transparency and accountability obligations on foreign firms than those imposed on domestic firms will invite scrutiny as a discriminatory measure – and so governments should be careful that any special scrutiny is properly justified. One question in this regard will be about a specific set of rules that are only triggered by size. If local companies are all likely to remain smaller than the threshold, there is the possibility of exploiting size triggers to disfavor foreign competitors. Furthermore, focusing only on the world's biggest Internet companies may or may not be justified because of their impact – but it is also important to remember that some of the most pernicious applications of AI might escape scrutiny if we limit our regulatory attention to a handful of enterprises.

Overall, today, decisions about people and machines are being made by machines. AI helps people file tax returns, it helps offer or deny loans, it matches individuals for dating, it makes investment decisions, sorts through job applications, and delivers search results. Given that AI is making decisions that affect people's lives, governments should insist on what we might call 'locally responsible AI.'

C AI AND TRADE LAW

Does trade law apply at all to AI? A skeptic might offer two arguments – the first textual and the second conceptual. First, the WTO agreements and the scheduled

⁹ On the inadequacy of legal transplants, see, e.g., P. Legrand, 'The Impossibility of 'Legal Transplants', *Maastricht Journal of European and Comparative Law* 4 (1997), 111–124; M. Siems, 'Malicious Legal Transplants,' *Legal Studies* 38 (2018), 103–119.

¹⁰ BBC, 'Volvo's Driverless Cars 'Confused' by Kangaroos,' *BBC News*, 27 June 2017, available at www.bbc.com/news/technology-40416606.

¹¹ B. Templeton, 'Another Alleged Tesla Autopilot Failure Raises Questions on Tesla Training System,' *Forbes*, 12 August 2019.

commitments of the WTO members that form an integral part of the treaties nowhere mention AI, and thus should not be interpreted to cover this new technology.¹² Applying trade law to this new sphere would violate the sound expectations of the parties. Second, AI is simply a method of doing something, the skeptic might assert, and the trade agreements focus on what is actually provided rather than the process used to provide it – a version of the process/product distinction elaborated for goods.¹³ After all, if trade law does not scrutinize whether a particular decision made by a company is made by an individual or a committee, then why should it pay attention to the decision-making process at all?

Can the WTO agreements apply to AI decision-making? Even if AI techniques were not widely used when the WTO agreements were negotiated, the General Agreement on Trade in Services (GATS)¹⁴ does not limit itself to the technologies in use in 1994. GATS proves relevant through three characteristics: First and most importantly, GATS focuses on measures regulating services without regard to the technologies by which those services are provided.¹⁵ Its first substantive sentence declares, '[t]his Agreement applies to measures by Members affecting trade in services.'¹⁶ Second, the GATS applies to technologies that may have not been on the minds of the negotiators.¹⁷ When China sought to deny that it had included *electronic* distribution of audiovisual material in its WTO commitments in the *China – Audiovisual Products* case, the WTO Appellate Body ruled decisively that it was indeed covered.¹⁸ As I have noted elsewhere, 'By subsuming an electronic version of the service within a services commitment and by interpreting treaty

¹² R. Zhang, 'Covered or Not Covered: That Is the Question,' WTO Working Paper No 11 (2015), at 14–17. 'A more far-reaching interpretation was that any service 'unforeseen' at the time of commitments could not be considered as covered by it, even if the definition in the CPC covered the 'unforeseen' service.' Ibid.

¹³ See G. Cook, 'Humpty Dumpty and the Illusion of 'Evolutionary Interpretation' in WTO Dispute Settlement,' in G. Abi-Saab et al. (eds), *Evolutionary Interpretation and International Law* (Oxford: Hart Publishing, 2019). See also Appellate Body Report, *China – Measures Affecting Trading Rights and Distribution Services for Certain Publications and Audiovisual Entertainment Products (China – Audiovisual Products)*, WT/DS363/AB/R, adopted 21 December 2009. In *China – Audiovisual Products* (at para. 396), the Appellate Body found that the terms in China's Schedule 'are sufficiently generic that what they apply to may change over time.'

¹⁴ General Agreement on Trade in Services, 1869 U.N.T.S. 183; 33 I.L.M. 1167 (1994), entered into force 1 January 1995 [hereinafter: GATS].

¹⁵ WTO, Work Programme on Electronic Commerce – Progress Report to the General Council, SL/74, 27 July 1999, at 4; see also WTO Panel Report, *United States – Measures Affecting the Cross-Border Supply of Gambling and Betting Services (US – Gambling)*, WT/DS285/R, adopted 10 November 2004 and WTO Appellate Body Report, *United States – Measures Affecting the Cross-Border Supply of Gambling and Betting Services (US – Gambling)*, WT/DS285/AB/R, adopted 7 April 2005.

¹⁶ Article I:1 GATS.

¹⁷ See generally S.-Y. Peng, 'Renegotiate the WTO 'Schedules of Commitments'?: Technological Development and Treaty Interpretation,' *Cornell International Law Journal* 45 (2012), 403–430.

¹⁸ *China – Audiovisual Products*, note 12, at para. 412.

commitments in a dynamic form, the treaty can take account of changing technologies.¹⁹ If a term is listed in a sufficiently generic fashion, it should be interpreted to cover activities that were not commercialized at the time of the listing.²⁰ Indeed, when it determined that electronic distribution of audiovisual recordings was covered by China's commitments, the Appellate Body observed that it was not necessary that such electronic distribution was feasible at the time when China acceded to the WTO.²¹ Thus, a generic commitment for market access for insurance decision-making under mode 1 (cross-border supply) should be read to cover AI-based decision-making as well. Third, as *China – Audiovisual Products* decision makes clear, the GATS applies to electronically mediated services – a fact essential to enable it to cover AI-powered services. Fourth, the GATS schedules explicitly include a variety of computer and related services in their ambit, with at least seventy-seven countries committing to liberalize trade in 'data processing services.'²² The end result is that when a government measure affects the ability of a foreign company to supply AI-based services into that country, GATS is applicable.

The second objection challenges the idea that AI can be reached by trade law on the ground that *how* a decision is made with respect to any service is not a proper subject of trade law. This is a version of the controversial process and production methods (PPMs) distinction from the realm of goods,²³ where an importing government may not be able to inquire into the process by which a product is produced, only evaluating its quality as it arrives at the border.²⁴ Steve Charnovitz divides PPMs into three types: (i) the how-produced standard; (ii) the government policy standard; and (iii) the producer characteristics standard.²⁵ Translating this into the domain of services, it

¹⁹ A. Chander, *The Electronic Silk Road: How the Web Binds the World Together in Commerce* (New Haven, CT: Yale University Press, 2013), at 156.

²⁰ *China – Audiovisual Products*, note 12, at para. 396; Cook, note 12.

²¹ *China – Audiovisual Products*, note 12, at para. 396.

²² R. Berry and M. Reisman, 'Policy Challenges of Cross-Border Cloud Computing,' *Journal of International Commerce and Economics* 4 (2012), 1–38, at 22 (noting that sixty countries have commitments on 'on-line information and/or data processing,' while seventy-six have commitments for data processing). My review with Usman Ahmed finds at least seventy-seven countries with 'CPC 843' commitments for data processing services, though some of these commitments may be narrower than all data processing services. See U. Ahmed and A. Chander, 'Information Goes Global: Protecting Privacy, Security, and the New Economy in a World of Cross-Border Data Flows,' *E15 Expert Group on the Digital Economy Think Piece* (2015). For all commitments and exceptions of the WTO members for computer and related services, see www.wto.org/english/tratop_e/serv_e/computer_e/computer_e.htm.

²³ The term 'processes and production methods' originated in the GATT agreement of 1979 on Technical Barriers to Trade and referred to product standards focused on the production method rather than product characteristics. For example, a law prohibiting the landing of fish caught using a driftnet entails a PPM. By contrast, a law prohibiting the sale of fish smaller than a prescribed size does not constitute a PPM.

²⁴ S. Charnovitz, 'The Law of Environmental 'PPMs' in the WTO: Debunking the Myth of Illegality,' *Yale Journal of International Law* 27 (2002), 59–110, at 64–65.

²⁵ *Ibid.*, at 67.

would mean that the importing government treats the foreign service provider differently because of (i) how it produced the service; (ii) the law governing that service in the exporting country; or (iii) the characteristics of the foreign service provider, respectively.

With respect to services, however, regulation often focuses on both, the provider and the process used, as it may be difficult to regulate the service directly. Licensing requirements, for example, often seek to assure that the individual performing the task has the relevant education, ethics, and experience to perform the service. In general, how a service is produced may be important to evaluate its quality – such as knowing whether an accountant or an engineer or a cybersecurity expert has followed the standard protocols.²⁶ Of course, much of the process used to provide the service could be inscribed in the service itself but it is often difficult to see the mark of that process directly. Thus, we often use other measures to evaluate the service – such as the prominence of the firm or the education of its employees or their use of a widely accepted method.²⁷ This is no less true with AI. Demands for explainability, for example, which have become common nowadays,²⁸ are often ultimately about a form of due process, including the ability to challenge a decision that one feels is unjust.

The following two sections explore two specific scenarios of the interaction between AI and international trade rules.

I Scenario One: Dr. AI

Imagine if a country bars unlicensed medical diagnosis, and interprets this requirement to bar all AI-based medical diagnosis, as there is no process for licensing an AI. What if a foreign company wishes to offer AI-based medical diagnosis into that country? Could it rely on the GATS commitments to liberalize trade in data-processing services to argue that the ban on AI medical diagnosis violated that country's WTO obligations?²⁹

²⁶ Chander, note 18, at 146 ('The measure of the quality of a service often involves not just the appraisal of the outcome but also the appraisal of the process by which the service was produced'); see also K. Nicolaidis and S. K. Schmidt, 'Mutual Recognition 'on Trial': The Long Road to Services Liberalization,' *Journal of European Public Policy* 14 (2007), 717–734, at 719 ('for services almost all regulations have to do with processes').

²⁷ The WTO secretariat's Trade in Services division similarly observes: 'Services are intangible and their supply often requires an interaction between the service supplier and consumer. This implies that consumers frequently cannot appreciate the quality of the service until they have consumed it.' See WTO Trade in Services Division, *Disciplines on Domestic Regulation Pursuant to GATS Article VI:4 – Background and Current State of Play*, June 2011, available at www.wto.org/english/tratop_e/serv_e/dom_reg_negs_bckgddoc_e.doc.

²⁸ See, e.g., U. Gasser and V. A. F. Almeida, 'A Layered Model for AI Governance,' *IEEE Internet Computing* 21 (2017), 58–62; A. Deeks, 'The Judicial Demand for Explainable Artificial Intelligence,' *Columbia Law Review* 119 (2019), 1829–1850.

²⁹ Of course, the aggrieved corporation could not seek to enforce GATS itself but could pressure its home state to do so.

The first step in making such a claim is to establish that the country had in fact committed to liberalize trade in such AI-based medical diagnosis services in the first instance. The market access and national treatment obligations, as we have said, rest on a nation's GATS schedule. This, in turn, raises difficult questions of classification. Suppose an AI performs the task of assessing whether a skin lesion is cancerous and does so via a smartphone app. Many but not all WTO members used the United Nations' Central Product Classification (CPC) in its provisional 1991 version³⁰ to schedule their liberalization commitments. The CPC has been revised numerous times since but these updates have not been reflected in the law of the WTO.³¹ Under the CPC scheme, human health services are classified as 'CPC 931,' with subdivisions for 'general' (93121) and 'specialized' (93122) health services, as well as other subdivisions. But perhaps the AI could be seen as a 'data processing service' (CPC 843) or a 'database service' (CPC 844) at the same time – after all the AI is an immense data processor and may rely on significant database functions? The GATS classification is designed to be exclusionary – that is, any given service should fall only under one category³² but it can be difficult to place many technologically powered services within the classification framework existing at the time of the WTO's founding.

The CPC itself provides interpretative rules, including two rules relevant here:

1. The category that provides the most specific description shall be preferred to categories providing a more general description; and
2. Composite services consisting of a combination of different services which cannot be classified by reference to (a) shall be classified as if they consisted of the service which gives them their essential character, in so far as this criterion is applicable.³³

³⁰ United Nations, Provisional Central Product Classification (CPC), UN Statistical Papers, Series M, No 77, Ver.1.1, E.91.XVII.7, 1991 [hereinafter: CPC]. The CPC is a classification based on the physical characteristics of goods or on the nature of the services rendered. Each type of good or service distinguished in the CPC is defined in such a way that it is normally produced by only one activity as defined in International Standard Industry Classification of all Economic Activities (ISIC). The CPC covers products that are an output of economic activities, including transportable goods, non-transportable goods and services.

³¹ See R. H. Weber and M. Burri, *Classification of Services in the Digital Economy* (Berlin: Springer, 2012), at 19.

³² A. D. Mitchell and N. Mishra, 'Data at the Docks: Modernizing International Trade Law for the Digital Economy,' *Vanderbilt Journal of Entertainment and Technology* 20 (2018), 1073–1134, at 1090 ('in a country's schedule, commitments on a service sector or subsector are exclusive; thus, a specific digital service (like the search engine services of Google) cannot be simultaneously classified under computer and related services (more specifically, data processing services), telecommunications services (online information and data processing services), and advertising services'). See also *US–Gambling*, note 14.

³³ CPC 1991, at 28.

If we assume that ‘medical diagnostic service’ is more specific than ‘data processing service,’ then an AI-based medical diagnostic service should properly be classified as a ‘medical diagnostic service.’ Thus, a commitment under CPC 843 for a data processing service is likely insufficient to grant a foreign AI medical diagnostic service provider market access and national treatment in that country without a relevant CPC 931 human health service commitment.

In *China–Electronic Payments*, the panel, however, questioned this approach, arguing that ‘the matter is not so obvious that we could confidently determine, without undertaking a detailed examination, [which service] is ‘more specific’ in relation to the services at issue.’³⁴ Yet, the panel’s preferred approach largely reached the same conclusion. The panel recognized ‘electronic payment services for payment card transactions’ as an ‘integrated service’ that included other services that could be provided independently.³⁵ The relevant classification in such cases would be the one describing the integrated service.³⁶

What if a country has left medical services unbound, but has bound data processing services for both market access and national treatment? Would a foreign AI medical diagnostic provider be able to benefit from that data processing commitment? It seems likely that it would only be able to claim them for providing data processing but not for the medical diagnostic service itself, which would have required a CPC 931 commitment.

The scheduling guidelines adopted by the WTO’s Council for Trade in Services in 2001 distinguish between a committed service and input services to that committed service.³⁷ The scheduling of a committed service does not imply that the input services are also equally committed when used for purposes other than the committed, composite service. It seems sensible, however, to assume that the input services are automatically committed when provided as an input into the committed service – that is, it should not be possible for a WTO member to specify that a foreign medical diagnostic provider (presuming that medical diagnostic services are committed) must use domestic AI. Otherwise, the commitment of the integrated service would be less meaningful because one could establish a variety of requirements for the inputs into that service that would greatly erode the commitment. Then if members specify medical diagnosis, they need not specify all the input services needed to supply a medical diagnosis. In our hypothetical case of ‘Dr. AI,’ if the data processing or database service is an input service to the AI-based medical diagnostic

³⁴ WTO Panel Report, *China – Certain Measures Affecting Electronic Payment Services (China–Electronic Payment Services)*, WT/DS413/R, adopted August 31, 2012, at para. 7.71.

³⁵ *Ibid.*, at paras. 7.55–7.62, and 7.188.

³⁶ *Ibid.*

³⁷ WTO, Guidelines for the Scheduling of Specific Commitments under the General Agreement on Trade in Services (GATS), S/L/92, 23 March 2001, at 25 (‘It is understood that market access and national treatment commitments apply only to the sectors or subsectors inscribed in the schedule. They do not imply a right for the supplier of a committed service to supply uncommitted services which are inputs to the committed service’).

service, then a commitment under CPC 931 for such a service would include the data processing or database service.

II Scenario Two: Claims Adjuster AI

Imagine a country that bans automated decision-making for insurance coverage decisions. This would go beyond the right to object to a decision made by an automated algorithm under the European Union's General Data Protection Regulation (GDPR).³⁸ Such a scenario would be reminiscent of the genetic engineering debate in trade law – where Europe rejected genetically modified food outright, while the United States insists on their safety.³⁹

Imagine also that domestic insurance providers are not technologically minded, while foreign competitors are more likely to use AI. So the burden of the rule largely falls on foreign providers. Assume that the country banning AI has made market access and national treatment commitments for the relevant insurance products under the Annex on Financial Services, but has limited those to mode 3 (commercial presence), as countries often are reluctant to allow for cross-border trade in financial services because of prudential regulation of financial institutions to ensure, among other things, their safety and soundness.

Might the foreign country of that foreign insurance provider with a domestic establishment have a claim? The foreign home might challenge the absolute bar as a violation of that importing state's market access commitments. A ban might be seen as a zero quota, and thus a numerical limitation on the number of providers – which will be a violation of the GATS market access obligation contained in Article XVI:2.⁴⁰

The foreign country might also argue that the ban violates the national treatment requirement by effectively preferring domestic insurance providers, which do not use AI for decisions. A central question in answering this question is whether the AI-based insurance service was 'like' the non-AI based insurance service. While guidance on the interpretation of 'likeness' when it comes to services is limited,⁴¹

³⁸ Regulation 2016/679 of the European Parliament and of the Council of 27 April 2016 on the Protection of Natural Persons with Regard to the Processing of Personal Data and on the Free Movement of Such Data, and Repealing Directive 95/46/EC (General Data Protection Regulation, GDPR), OJ L [2016] 119/1 [hereinafter: General Data Protection Regulation or GDPR], at Article 22(1): data subjects 'have the right not to be subject to a decision based solely on automated processing, including profiling, which produces legal effects concerning him or her or similarly significantly affects him or her.'

³⁹ See WTO Panel Report, *European Communities – Measures Affecting the Approval and Marketing of Biotech Products (EC–Biotech)*, WT/DS921/R, adopted 21 November 2006; see also M. A. Pollack and G. C. Shaffer, *When Cooperation Fails: The International Law and Politics of Genetically Modified Foods* (Oxford: Oxford University Press, 2009).

⁴⁰ WTO Panel Report, *US – Gambling*, note 14, at para. 239.

⁴¹ The US argued that online gambling was unlike real world gambling because 'in virtual casinos, the result is generated by a software algorithm' rather than physical movement. The 'online casino is an illusion – a 'virtual reality' environment in which outcomes are controlled

the Appellate Body has indicated that the ‘fundamental purpose’ of the likeness comparison is ‘to assess whether and to what extent the services and service suppliers at issue are in a competitive relationship.’⁴²

If a tribunal concludes that the AI ban violates either market access or national treatment commitments, the importing nation will argue that the ban is justified by considerations of privacy, public order, or even public morals (with respect to the latter, the argument would be that having such important decisions made as insurance denial about someone by an AI would be an affront to human dignity). Article XIV of the GATS permits a derogation that is ‘necessary to protect public morals or to maintain public order’⁴³ but the ‘public order exception may be invoked only where a genuine and sufficiently serious threat is posed to one of the fundamental interests of society.’⁴⁴ One focal point of the analysis will be whether the ban is *necessary* to protect public order. The exporting nation might argue that an alternative WTO-consistent that achieves the same ends is reasonably available, and thus an outright ban is not necessary.⁴⁵ It might for instance point to the German approach as such an alternative: Germany explicitly recognizes automated decision-making for insurance decisions but requires the insurance company to offer human review for any negative decisions.⁴⁶

In summing up, even if existing trade law does have mechanisms to reduce protectionist barriers to trade in AI, there remains substantial room for disagreement over whether any particular rule that burdens trade in AI can be justified. The examples above point to some of the debates and critical questions, such as: Is AI medical diagnosis ‘like’ human medical diagnosis? Can an AI-based insurer be banned on the grounds that it is likely to be biased or opaque? The rules as they

by a computer rather than by the laws of the physical world,’ the US insisted (see WTO Panel Report, *US–Gambling*, note 14). Because of the exercise of judicial economy, however, the dispute settlement body did not reach the issue of national treatment, and thus, ‘there is up until today no jurisprudence on how such characteristics on the method of supply should be evaluated’ with respect to services. See N. Diebold, *Non-discrimination in International Trade in Services: ‘Likeness’ in WTO/GATS* (Cambridge: Cambridge University Press, 2010), at 252.

⁴² WTO Appellate Body Report, *Argentina – Measures Relating to Trade in Goods and Services*, WT/DS453/R, adopted 9 May 2016, at paras. 6.33–6.34.

⁴³ Article XIV(a) and Article XIV(b) GATS.

⁴⁴ Article XIV GATS.

⁴⁵ WTO Panel Report, *US–Gambling*, note 14, at paras. 306–307: ‘A comparison between the challenged measure and possible alternatives should then be undertaken, and the results of such comparison should be considered in the light of the importance of the interests at issue. It is on the basis of this ‘weighing and balancing’ and comparison of measures, taking into account the interests or values at stake, that a panel determines whether a measure is ‘necessary’ or, alternatively, whether another, WTO-consistent measure is ‘reasonably available.’ The exporting country can also argue that the challenged measure is arbitrary or is an unjustifiable discrimination between countries where like conditions prevail, or a disguised restriction on trade in services, as specified in the chapeau of Article XIV GATS.

⁴⁶ G. Malgieri, ‘Automated Decision-Making in the EU Member States: The Right to Explanation and Other ‘Suitable Safeguards’ in the National Legislation,’ *Computer Law and Security Review* 35 (2019), 1–26, at 7–8.

stand do not give clear answers to such questions. Internationally agreed frameworks for responsible AI might offer a process to protect national regulatory goals while enabling trade in AI.

D CONCLUSION

Governments across the world are struggling to keep up with technology. The rise of AI decision-making, in everything from cars to media to business processes, challenges regulatory capacity. Governments must regulate AI in order to further traditional regulatory goals, such as consumer protection, privacy, and law enforcement. Governments can, however, craft or enforce AI rules that disfavor foreign enterprises. The regulation of AI should not be used to create yet another behind-the-border trade barrier.