Singapore Journal of Legal Studies
[2022] 46–74

B2B ARTIFICIAL INTELLIGENCE TRANSACTIONS: A FRAMEWORK FOR ASSESSING COMMERCIAL LIABILITY

Ernest **Lim***

Business to business ("B2B") artificial intelligence ("AI") transactions raise challenging private law liability issues because of the distinctive nature of AI systems and particularly the new relational dynamics between AI solutions providers and procurers. This article advances a three-stage framework comprising data management, system development and implementation, and external threat management. The purpose is to unpack AI design and development processes involving the relational dynamics of providers and procurers in order to understand the parties' respective responsibilities. Applying this framework to English commercial law, this article analyses the potential liability of AI solutions providers and procurers under the *Supply of Goods and Services Act* and the *Sale of Goods Act*. The assumption that only AI solutions providers will be subject to liability, or that no party will be liable due to the "autonomous" nature of AI systems, is rejected.

I. Introduction

Businesses increasingly deploy artificial intelligence ("AI") systems in diverse areas such as finance,[1] healthcare,[2] taxation,[3] sales and marketing,[4] production and manufacturing,[5] and risk management.[6] In the absence of in-house expertise, many

* Professor, Faculty of Law, National University of Singapore. My greatest debt of gratitude goes to Iris Chiu for her extensive and insightful input, without which this article would not have materialised. I also wish to express my appreciation to Michael Bridge, Christian Witting and Tan Zhong Xing for their substantive and constructive comments on an earlier draft of this paper. I would also like to thank Satya Kumar, Lee Kay Han and Jonathan Kow for their research assistance. The usual disclaimers apply.

1 Christian L Dunis *et al*, eds, *Artificial Intelligence in Financial Markets: Cutting-Edge Applications for Risk Management, Portfolio Optimization and Economics* (London: Palgrave Macmillan, 2016).

2 Bernard Nordlinger, Cédric Villani & Daniela Rus, eds, *Healthcare and Artificial Intelligence* (New York: Springer, 2020).

3 PwC, "How Tax is leveraging AI—Including Machine Learning—In 2019" (2019), online: <https://www.pwc.com/cb/en/services/pdf/how-tax-leveraging-ai-machine-learning-2019.pdf>.

4 Thomas Davenport *et al*, "How Artificial Intelligence Will Change the Future of Marketing" (2020) 48 J Acad Marketing Sci 24; Peter Gentsch, *AI in Marketing, Sales and Service* (London: Palgrave Macmillan, 2019).

5 Ray Y Zhong *et al*, "Intelligent Manufacturing in the Context of Industry 4.0: A Review" (2017) 3 *Engineering* 616; Matthias Finger *et al*, "Regulation for Artificial Intelligence and Robotics in Transportation, Logistics and Supply Chain Management" (2018) 20 Network Industries Q 1.

6 Saqib Aziz & Michael Dowling, "AI and Machine Learning for Risk Management" in Theo Lynn *et al*, eds, *Disrupting Finance: FinTech and Strategy in the 21st Century* (London: Palgrave, 2019); Federation of European Risk Management Associations ("FERMA"), "Artificial Intelligence Applied to Risk Management" (2019) FERMA Paper.

companies engage third parties to supply them with AI solutions. But these AI systems may perform in an unexpected or defective manner or cause losses to companies that have procured them (the "procurers"[7]). Thus, a key issue is whether AI solutions providers could be liable under English commercial law, specifically for breaching the *Supply of Goods and Services Act 1982*[8] ("*SGSA*") and the *Sale of Goods Act 1979*[9] ("*SGA*").

However, liability issues arising from business to business ("B2B") AI dealings are distinct from those in non-AI contexts (such as the sale of ordinary computer systems) because of two separate but interrelated features in B2B AI transactions: the nature of AI systems, and the relational dynamics between AI solutions providers and procurers.

Regarding the first feature, AI systems that deploy machine learning are unpredictable and opaque, as the design and functions of machine learning algorithms[10] differ from those in ordinary computer programming, as the former do not operate linearly using distinct symbols unlike the latter. Instead, machine learning algorithms process data in unpredictable ways as the system adapts to the evolving real environment, known as the "black box" problem.[11] Thus, if AI solutions providers lack complete control or knowledge[12] of the AI machine learning systems that they supply, one concern is that it may be difficult to impose liability on them for defective AI systems. But it is contended that it does not follow that providers know nothing about the data that has been selected and inputted into the AI systems.

This in turn leads to the second feature. Although the processing of data by machine learning algorithms and the bases on which the outcomes generated by AI systems are opaque and unpredictable, both AI solutions providers and procurers could play a role in constructing the AI system, such as by selecting and inputting the relevant data into the system. By contrast, ordinary computer systems are generally in a "ready-made" state that can be immediately used by the procurers who are not involved in the construction of the computer systems.[13] Thus, in a B2B AI context, it is important to understand the roles played by providers and procurers and the particular stage of the development and implementation of the AI system because the responsibilities of providers and procurers will differ. However, there appears to be no critical analysis of this (second) feature, particularly its

---

7 "Procurer" includes (1) the buyer under the *Sale of Goods Act 1979* (UK), c 54; (2) the party to whom the supplier agreed to supply the service under the *Supply of Goods and Services Act 1982* (UK), c 29; and (3) the party who engaged another party to use AI systems for certain functions including trading shares, an example of which is the case of *Tyndaris Sam v MMWWVWM Ltd* [2020] EWHC 778 (Comm) [*Tyndaris*] analysed in Part III(B).

8 (UK), c 29.

9 (UK), c 54.

10 Woodrow Barfield, ed, *The Cambridge Handbook of the Law of Algorithms* (Cambridge: Cambridge University Press, 2020).

11 *Eg*, Davide Castelvecchi, "Can We Open the Black Box of AI?" (5 Oct 2016), *Nature*, online: <https://www.nature.com/news/can-we-open-the-black-box-of-ai-1.20731>.

12 *Eg*, Mark Coeckelbergh, "Artificial Intelligence, Responsibility Attribution, and a Relational Justification of Explainability" (2019) Sci Engineering Ethics.

13 See *eg*, cases cited at *infra* notes 81, 84, 85.

implications for commercial law (although the first feature has been covered in the literature).

This article advances a three-stage framework comprising (1) data management, (2) system development and implementation, and (3) external threat management. The purpose is to unpack AI design and development processes involving the relational dynamics of AI solutions providers and procurers in order to understand the parties' respective responsibilities and to assess their potential liabilities under commercial law. At stage one, data management may contain risks of inaccurate, incomplete and unrepresentative data inputted by AI solutions providers or procurers, or both. Thus, procurers and providers should keep data description and data tracking records. At stage two, system development and implementation may contain problems including insufficient AI model training and testing, and usage error. The importance of ensuring that the relevant party has the experience and expertise to test and operate the AI system is illustrated with a stock trading example. Finally, at the last stage, managing external threats includes providers' and procurers' corresponding responsibilities in pre-empting and addressing cybersecurity attacks.

Applying the above framework to commercial law, the following key arguments are made. First, a stronger case can be made for B2B AI transactions to be characterised as contracts for the supply of services under the *SGSA*. Further, it is argued that AI software is unlikely to count as 'goods' under these legislation. Nevertheless, the article explains why the *SGA* is still relevant and the analysis will include both statutes.

Second, under the *SGSA*, whether the provider has supplied the AI system with reasonable care and skill depends on at which stage of the three-stage framework the care and skill has been exercised, and by which party. The framework enables the nature and scope of responsibilities that have been assumed by providers and procurers to be delineated. After all, procurers could be responsible, for example, for selecting or feeding erroneous or insufficient data into the AI system. It is also argued that providers should not be permitted to evade liability by arguing for a lack of causation due to the inherent unpredictability of machine learning.

Third, under the *SGA*, whether and the extent to which liability will be imposed on providers for the AI system's unsatisfactory quality will depend on whether providers explicitly and clearly inform procurers that inaccurate, incomplete or non-representative data will make the AI system's quality unsatisfactory; it will also depend on the clarity and precision by which providers circumscribe their role at the testing and training stage. It is also shown that whether providers comply with the fitness for purpose requirement depends on whether it is reasonable for procurers to rely on providers' skill and judgement at the different stages of the framework.

Finally, in view of the framework, contractual clauses seeking to exclude providers' liability under the *SGSA* and *SGA* could be construed on the whole as unreasonable under the *Unfair Contract Terms Act 1977*[14] ("*UCTA*"). The validity of duty-defining clauses is also assessed.

---

[14] (UK), c 50.

II. AI Systems in B2B Contexts

To understand why the liability issues arising from B2B AI transactions are different from B2B non-AI transactions, the question of what makes AI systems different from non-AI ones has to be addressed.

A. *What is AI?*

There are various definitions of AI. Fundamentally, AI are systems designed to reason and act like intelligent and rational human beings to attain specified objectives.[15] Examples of AI include autonomous vehicles, home management systems (the internet of things), predictive analytics including algorithms deciding on loan applications and job applicant hires.[16] Automation and autonomy are distinguished. Automation refers to pre-programmed systems (including advanced missiles, drones or autopilots, high-frequency trading software,[17] and industrial robots). Autonomous refers to systems incorporating and processing data, learning from patterns and making decisions, mostly for predictions requiring limited human intervention. Such AI systems may process data using various techniques including machine learning.[18] Depending on what "learning" paradigm is selected in machine learning, algorithms coded in AI systems rely on fed data to generalise, predict and generate outcomes. Machine learning includes or can be applied to: natural language processing,[19] letting human language expressions be directly engaged with instead of translated into code; decision trees[20] letting information analysis and processing pathways be organised with statistical and consequential logic; or artificial neural networks[21] simulating human brains' associations and organising data in statistical but non-linear manners (collectively, "machine learning routes"). Machine learning takes three forms—supervised learning, unsupervised learning, and reinforcement learning.[22] Supervised learning requires entering labelled data in a controlled environment to generate predicted output. Unsupervised learning

---

15 European Commission, High-Level Expert Group on Artificial Intelligence, *A Definition of AI: Main Capabilities and Scientific Disciplines* (Brussels: EC, 2018) at 1.

16 James Maguire, "12 Examples of Artificial Intelligence: AI Powers Business" (13 Sep 2019), *Datamation*, online: <https://www.datamation.com/artificial-intelligence/12-examples-of-artificial-intelligence-ai-powers-business/>.

17 *Eg* the defendant's trading software in *Quoine Pte Ltd v B2C2 Ltd* [2020] SGCA(I) 02. But not the AI machine learning software deployed by Tyndaris: see Part III(B).

18 Thomas P Trappenberg, *Fundamentals of Machine Learning* (Oxford: Oxford University Press, 2019).

19 Alexander Clark, Chris Fox & Shalom Lappin, eds, *The Handbook of Computational Linguistics and Natural Language Processing* (New Jersey: Wiley-Blackwell, 2012).

20 Lior Rokach & Oded Maimon, *Data Mining with Decision Trees: Theory and Applications*, 2d ed (Singapore: World Scientific, 2008), online: <http://repository.fue.edu.eg/xmlui/bitstream/handle/123456789/3632/9792.pdf>.

21 *Eg* Mohammed Abbas Kadhim, Afshar Alam & Harleen Kaur, "A Multi-Intelligent Agent for Knowledge Discovery in Database (MIAKDD): Cooperative Approach with Domain Expert for Rules Extraction" in De-Shuang Huang *et al*, eds, *Intelligent Computing Methodologies* (Heidelberg: Springer, 2014), 602.

22 Ian Goodfellow, Yoshua Bengio & Aaron Courville, *Deep Learning* (Massachusetts: MIT Press, 2016); Isha Salian, "SuperVize Me: What's the Difference Between Supervised, Unsupervised,

relies also on data—albeit not pre-labelled (unlike supervised learning)—to draw new inferences from the data set. Semi-supervised learning combines both. In reinforcement learning, AI has discretion to decide in complex, uncertain environments based on data available in the environment; it decides based on trial and error, and on feedback from its past experiences.

B. *How is AI Different and What are the Implications?*

For example, in sales and marketing, Intel uses AI, specifically machine learning.[23] Before using AI, Intel's sales and marketing analysts manually searched companies, identifying potential sales leads. With machine learning, Intel discovered new, better leads faster and more accurately. Intel developed an in-house AI system identifying new markets and customers using machine learning (specifically supervised and semi-supervised learning and natural language processing models) to segment customers. Intel fed millions of web-sourced textual data (including thousands of company sites appearing in Wikipedia) into a third-party-developed neural network text classification model, with a pre-trained, Google-developed multi-lingual language model. Intel labelled the data following two categories: industries (retail, transportation, education, healthcare, communications etc); and roles (service providers, retailers, manufacturers etc). For unlabelled companies, Intel deploys semi-supervised learning which lets the system freely determine the label, drawing from Intel's internal data from existing client relationships.

The Intel example shows how AI systems are different from non-AI ones such as ordinary computer systems in two crucial respects: the black box nature of AI systems, and the roles played by different parties in the data aspects of AI systems.

As to the first respect, Intel's AI machine learning system depends on artificial neural networks, containing neurons operating on inputted data.[24] Artificial neural networks are a series of algorithms containing neurons (nodes) connected by data-carrying links (*ie*, input).[25] The neurons operate from data received via the links. The neurons assign weights to the inputs based on various statistical and non-linear paradigms to arrive at a prediction. Thus, machine learning is dynamic and adaptive—often termed the "black box" problem.[26] By contrast, in ordinary programming, the algorithms generally operate linearly using distinct symbols, and thus programmers can explain the algorithms' precise workings, and third parties can check the algorithms' accuracy. Accordingly, in view of the black box nature of AI machine learning, an issue is whether and if so how responsibility and liability can be ascribed, and to whom, should the AI system fail and cause losses to the

---

Semi-Supervised and Reinforcement Learning?" (2 Aug 2018), *The Official NVIDIA Blog*, online: <https://blogs.nvidia.com/blog/2018/08/02/supervised-unsupervised-learning/>.

[23] Itay Lieder *et al*, "Learning a faceted customer segmentation for discovering new business opportunities at Intel", online: <https://arxiv.org/pdf/1912.00778.pdf>.

[24] Anukrati Mehta, "A Comprehensive Guide to Types of Neural Networks" (25 Jan 2019), *DigitalVidya*, online: <https://www.digitalvidya.com/blog/types-of-neural-networks/>.

[25] Larry Hardesty, "Explained: Neural Networks" (14 April 2017), *MIT News Office*.

[26] *Eg*, Yavar Bathaee, "The Artificial Intelligence Black Box and the Failure of Intent and Causation" (2018) 31:2 Harvard J L Tech 890.

party who has procured it. Despite the black box nature of AI, it does not follow that no responsibility can be attributed to AI solutions providers.

This leads us to the second respect in which AI systems differ from non-AI ones. Unlike non-AI systems (such as ordinary computer systems) whose design and functioning generally do not depend on the data that has been selected and inputted into them by different parties, AI machine learning systems are different. The internal architecture of artificial neural networks changes with the type and extent of data inputted by different parties, and hence the AI system may be unpredictable in terms of its performance.[27] Thus, the data used in AI systems are critical. Some of the data are drawn from internal sources and some from external ones. Companies in the roles of AI solutions providers and procurers may be responsible to different extents in inputting data to AI systems under development and implementation. There is a risk that the data may be inaccurate, incomplete and/or unrepresentative. In the example above, Intel is primarily responsible for selecting and feeding the data, but the artificial neural network text classification system into which Intel feeds the data also contains and relies on prior data that have been inputted by other parties. Further, how data are being labelled by AI solutions providers or procurers is a crucial issue, as erroneous labelling can result in the AI system deviating from intended outcomes. Consequently, a critical issue is who should be responsible and for what depends on the precise role played by the party in relation to how the data are being handled.

An illustration of the dual role played by providers and procurers in the design and functioning of AI systems, specifically data handling, can be found in the AI solutions designed by auditing companies such as Deloitte (the AI solutions provider). Deloitte analyses its corporate customer's (*ie*, procurer's) employment tax obligations in order to promote more effective tax compliance.[28] Deloitte initially feeds the AI model its own data including dictionary, tax laws and regulations, and training data. The company (with Deloitte's help) then locates, extracts and analyses data from the company's general ledger, payroll, and accounts payable systems. All relevant data are then labelled according to different employment-related expenditures. The AI system is then trained with different sets of scenarios and questions to see whether it provides correct tax compliance answers. In another case study, Deloitte used machine learning to create an AI system extracting clauses in contracts, deeds, and trust documents.[29] Corporate customers have used that system to review trust documents to determine appropriate tax regimes for trust arrangements.

The above examples (and there are many more[30]) demonstrate the importance of data management and training in AI systems. These have important implications for the legal analysis of parties' responsibilities and liabilities in B2B contexts. Part III explains the three-stage framework for AI development and implementation to show

---

27 Charu C Aggarwal, *Neural Networks and Deep Learning* (New York: Springer, 2018).

28 Deloitte, "Artificial Intelligence—Entering the World of Tax" (2019), *Deloitte* at 5, online: <https://www2.deloitte.com/global/en/pages/tax/articles/artificial-intelligence-in-tax.html>.

29 *Ibid*.

30 *Eg*, Kanetix, "Kanetix Ltd. Partners with Integrate.ai to Become First InsurTech Company to Use Machine Learning to Optimize Consumer Experience" (26 Feb 2018), *Cision*; deepsense.ai, "Credit card cross-selling", online: <https://deepsense.ai/portfolio-item/credit-card-cross-selling/>.

the technical and operational risks pertaining to data that arise at different stages of AI systems. Part IV then explains how the framework can illuminate the liability issues that arise in commercial law.

III. The Three-Stage Framework

The three-stage framework examines three key aspects of AI design and development: first, data management; second, system development and implementation; and finally, external threat management.[31] Data management encompasses which, and how, data are selected, labelled and fed into the AI system. The risks at this stage include incomplete, inaccurate and unrepresentative data. System development and implementation encompasses how the AI model is tested, trained, validated and deployed. The main risks at this stage are insufficient training and testing, and usage error. External threat management pertains to the risks of hacking and data breach. This three-stage framework fleshes out the complex relational dynamics between the AI solutions provider and procurer, thereby allowing us to feed into the legal analysis of potential liability in cases involving procurers (*ie*, claimants) and AI solutions providers (*ie*, defendants) in B2B contexts.

A. *Data Management—Incomplete, Inaccurate or Non-representative Data*

AI systems' reliability depends on the quality and quantity of data selected, fed into and processed by the system.[32] For instance, machine learning in AI systems can be used to detect tumours in connection with cancer treatments. Such AI systems process fed data, including patients' records, tumour images from many patients, public sources, different diagnosis and treatment methods, and AI solutions providers' proprietary medical sources.[33] There are two data suppliers: providers and procurers' data scientists. Beyond tumour detection, AI using machine learning detects chances of stroke in patients by relying on data such as the paraesthesia of the arms and legs.[34] However, AI systems can wrongly diagnose by stating that cancer patients are cancer-free (false negatives) or that cancer-free patients have

---

[31] This framework is a significant revision and expansion of Benjamin Cheatham, Kia Javanmardian & Hamid Samandari, "Confronting the Risks of Artificial Intelligence" (26 April 2019), *McKinsey Quarterly* at 5 and partially draws upon Lucy Ellen Lwakatare *et al*, "A Taxonomy of Software Engineering Challenges for Machine Learning Systems: An Empirical Investigation" in Philippe Kruchten, Steven Fraser & François Coallier, eds, *Agile Processes in Software Engineering and Extreme Programming* (New York: Springer Open, 2019, Lecture Notes in Business Information Processing vol 355).

[32] Mary H Stanfill & David T Marc, "Health Information Management: Implications of Artificial Intelligence on Healthcare Data and Information Management" (2019) 28 Yearbook Med Informatics 56.

[33] Arjun Panesar, *Machine Learning and AI for Healthcare: Big Data for Improved Health Outcomes* (New York: Apress, 2019).

[34] Casey Ross & Ike Swetlitz, "IBM pitched its Watson supercomputer as a revolution in cancer care. It's nowhere close" (5 Sep 2017), *STAT*, online: <https://www.statnews.com/2017/09/05/watson-ibm-cancer/>.

cancer (false positives). Such errors can be attributed to inaccurate or incomplete data from, for instance, mislabelled medical scans.[35] Such an error occurred when IBM fed non-representative data into its Watson system. Data of a small number of synthetic, nonreal cases were included, owing to limited data from oncologists' real cases.[36] Regarding the sales and marketing of AI systems, ensuring the accuracy of marketing data is a challenge in view of the vast amount of data produced from consumer interactions with a myriad of brands and companies.[37] Moreover, customer data fed by a company into its AI system can contain spelling errors, for instance.[38]

It is important to determine at which stage the data was selected and fed, and by whom. AI solutions providers could be responsible for feeding initial data during AI system testing and design, and procurers could be jointly responsible for data input at training, and solely responsible for post-deployment data input. Others may also curate and feed data. For example, the Deloitte-designed AI tax system contains Deloitte-fed data including relevant laws and regulations at the initial design stage.[39] But the procurer (the corporate customer) selects and feeds ledger and operational data.[40] If Deloitte's system contains incomplete data because it omitted material tax regulations, therefore causing the AI system to recommend an erroneous tax structuring, then Deloitte's responsibility at the data management stage may be crucial for ascertaining causation and liability.

To address the problem of inaccurate, incomplete or non-representative data, it is important to know the data's origins, data collection and curation methods, and the precautions taken to minimise inaccuracy.[41] Measures clearly tracing data movement, and whether and how data was altered therein, should be implemented. Procurers and AI solutions providers should keep coherent joint records of data movement to speedily identify which data were fed into the AI system, from where, and their reliability.[42] A clear process should also track how and when data are updated or corrected, particularly given regulatory developments. Developing such

---

35 David Gorski, "IBM's Watson Versus Cancer: Hype Meets Reality" (11 Sep 2017), *Science-Based Medicine*, online: <https://sciencebasedmedicine.org/ibm-watson-versus-cancer-hype-meets-reality/>.

36 Eric J Topol, "High-Performance Medicine: The Convergence of Human and Artificial Intelligence" (2019) 25 Nature Med 44.

37 Jim Sterne, *Artificial Intelligence for Marketing: Practical Applications* (New Jersey: Wiley, 2017).

38 *Ibid*.

39 Deloitte, *supra* note 28; PwC, *supra* note 3.

40 Deloitte, *ibid* at 5.

41 Singapore, Info-communications Media Authority ("IMDA") & Personal Data Protection Commission, *Model Artificial Intelligence Governance Framework*, 2d ed (Singapore: 2020); Thomas C Redman, "If Your Data is Bad, Your Machine Learning Tools Are Useless" (2 April 2018), *Harvard Business Review*, online: <https://hbr.org/2018/04/if-your-data-is-bad-your-machine-learning-tools-are-useless>.

42 The sources of data may be derived from not only the procurer and provider, but also from third parties through contractual arrangements. For example, the data concerning the tax laws and regulations that are fed by Deloitte into the AI tax system may be drawn from public or private entities with which Deloitte has entered into contracts in order to use and disseminate the data. In this respect, the party that fed the data, *ie* Deloitte in this example should check the accuracy and completeness of the data. The provider (Deloitte) should not be allowed to unilaterally exclude liability for materially incomplete or materially inaccurate data simply because the data was prepared by third parties (even if the procurer has contractually agreed to such disclaimer) unless the exemption clause satisfies the reasonableness test under the *UCTA*, analysed in Part IV(F).

records helps determine responsibility and eventually legal liability in B2B contexts. In addition to data movement records, it is suggested procurers and AI solutions providers should maintain joint data description records providing information on the content and quality of datasets. The responsibility for data quality may give rise to issues as to who should be responsible for data quality in a particular respect. In sum, data records should clearly and comprehensively detail data creation, composition, collection, and distribution, identifying parties' responsibilities. These have implications for the analysis of liability between procurers and providers, as examined in Part IV.

B. *System Development and Implementation—Insufficient Training and Testing as well as Usage Error*

At the system development and implementation stage, the AI model must be tested and trained.[43] Algorithms coded in AI systems rely on fed data to generalise, predict, and generate outcomes. This stage also comprises human interventions to the extent humans must operate AI systems such as correcting mislabelled data; overriding plainly erroneous system decisions; and feeding new or updated data into the system.

Two specific risks are connected to system development and implementation: first, insufficient training and testing; and second, usage error. To give a concrete explanation of the former using artificial neural networks,[44] consider the lawsuit by Tyndaris (an investment company) against VWM (an investment fund).[45] Under the contract between the claimant (Tyndaris) and the defendant (VWM), the claimant made investment decisions based on a supercomputer that is run on machine learning. The algorithms in the systems were supposed to gather and analyse data from real-time public sources (including news feeds and social media) to determine investors' sentiments, and then make trading decisions. However, Tyndaris' investment decisions resulted in VWM's suffering approximately USD22 million in investment losses. VWM instructed Tyndaris to stop trading on its behalf. Tyndaris then sued VWM for unpaid fees of USD3 million. VWM counterclaimed for the losses, citing its reliance on Tyndaris' misrepresentations concerning the supercomputer's capabilities. Key issues raised in the

---

[43] Ron Schmelzer, "How Do You Test AI Systems?" (3 Jan 2020), *Forbes*; Sanjay Nambiar & Prashanth Davey, "Testing of AI/ML based systems: The related challenges and effective techniques to overcome them", *WIPRO*, online: <https://www.wipro.com/holmes/testing-of-ai-ml-based-systems/>. The dataset is divided into mainly training (where the algorithm learns from the data to create patterns and make predictions), validation (where the model is evaluated with data to correct any prediction errors), and testing (where data will be used to provide a final evaluation of the AI model's performance post-training and -validation).

[44] Rajneesh Malviya *et al*, "The Human Touch in AI-Aided Trading" (2019), *Infosys*, online: <https://www.infosys.com/iki/insights/ai-aided-trading.html#:~:text=A%20paradigm%20shift%20is%20taking,and%20manage%20AI%2Daided%20trading.>.

[45] *Tyndaris*, *supra* note 7; Minesh Tanna & William Dunning, "AI-Powered Investments: Who (If Anyone) Is Liable When It Goes Wrong? Tyndaris v VWM" (2020), *Simmons & Simmons*, online: <https://www.simmons-simmons.com/en/publications/ck2xifd2ddmrq0b48u46j2nns/ai-powered-investments-who-if-anyone-is-liable-when-it-goes-wrong-tyndaris-v-vwm>.

documents filed with the High Court include those related to system development and implementation such as whether Tyndaris had adequate training and expertise to operate the supercomputer; whether Tyndaris properly tested the supercomputer; and whether Tyndaris operated the supercomputer in a manner consistent with its contractual duties to VWM.

For example, adequately training and properly testing an AI trading model using artificial neural networks should include the following steps.[46] First, publicly acquiring share price data followed by real-time analyses of them, and organising voluminous data based on certain criteria. Second, developing the system, including selecting appropriate artificial neural network types for deployment,[47] then feeding the data into the neural networks. Finally, the AI system's ability to generate outcomes predicting market price movements must be tested, thereby ensuring sound investment decisions. Based on the satisfactory (or otherwise) predictions or recommendations the artificial neural networks generate, the AI system may need retraining, for example by relabelling data or feeding corrected or updated data. Thus, even with proper initial data management, AI solutions providers could be responsible for inadequate or improper AI model testing. Moreover, despite the inherent unpredictability in artificial neural networks' predictive processes, the AI solutions provider must be responsible for model testing and training. Sub-optimal testing and training will exacerbate inherent unpredictabilities (or in clear cases cause the AI system not to perform in anything like the way it was designed to do).

It is pertinent to interrogate the procurer's and AI solutions provider's respective responsibilities in the system development and implementation stage. Each of their responsibilities is likely to be mapped onto an intricate web of relational dynamics that may be contractually agreed upon, and informally developed over a course of conduct. Further, jurisprudence must develop the normative aspect of who should be responsible. AI solutions providers may generally be better-placed to choose particular machine learning routes for AI system development, and should be better-placed to be responsible for observing the effects of system tests and identifying necessary corrections.

If the AI system was properly trained, tested and validated, another problem relates to usage error, *ie*, how the procurer used it. The procurer could be responsible for exercising expert judgment to override the AI system's recommendations if the recommendations are not acceptable according to expert judgment. For example, in healthcare, medical professionals must exercise judgment to override or reject AI system recommendations that deviate from well-established medical practices and procedures. Further, procurers contravening agreed or established AI system usage procedures should bear responsibility, not the AI solutions providers.

---

[46] Vivek Palaniappan, "Neural Networks to Predict the Market" (23 Sep 2018), *Towards Data Science*, online <https://towardsdatascience.com/neural-networks-to-predict-the-market-c4861b649371>.

[47] Kishnan Maladkar, "6 Types of Artificial Neural Networks Currently Being Used in Machine Learning" (15 Jan 2018), *AIM*, online: <https://analyticsindiamag.com/6-types-of-artificial-neural-networks-currently-being-used-in-todays-technology/>; Mehta, *supra* note 24.

C. *Managing External Threats*

Finally, external threats to AI systems including cyber hacking must be managed. Fraudsters can hack into and exploit sensitive or confidential data in AI systems. This is particularly serious with regard to patients' data in AI medical systems and taxpayers' information in AI tax systems.[48] AI systems' algorithms can also be exploited or manipulated, affecting outcomes generated for users. Is it the procurers' or AI solutions providers' responsibility to install cybersecurity protections for AI systems? Even if the answer to this question is clear, under contractual arrangements, the other party possibly has a duty to ensure the AI system is protected from external threats within their capability.

D. *The Three-Stage Framework: Concluding Considerations*

In light of the three distinct stages in the framework, both AI solutions providers and procurers should identify with care and precision at which stage and for what each of the parties is responsible in the contract. Thus, for example, it is possible for the provider and procurer to contractually agree that only the latter will be responsible for the deployment and maintenance of the AI system, *ie*, the second and third stages of the framework. Nevertheless, parties should pay attention to how their respective responsibilities are demarcated not only between the three stages, but also within each of the stages. This is because, for example, the second stage consists of not only operation/deployment of the AI system after it has been delivered to the procurer, the responsibility of which usually falls on the latter, but also training and testing the AI model, the responsibility of which usually lies with the provider. Moreover, one should pay attention to the fact that the responsibility of one party at a particular stage may depend on the other party discharging its responsibility. For example, the provider's correct and complete testing and training of the AI model (the second stage) may depend on the updated or corrected data supplied by the procurer (which data are different from those that are initially furnished by the procurer during first stage, *ie*, data management). Further, one should not assume that just because it has been mutually agreed that responsibility for a particular stage (such as stage three) lies with one of the parties (the procurer), the other party (the provider) is completely exempted from any responsibility. For example, it should not be assumed the procurer will be automatically responsible for the damages arising from not installing cybersecurity protection (the third stage), if doing so would not have prevented the cyber hacking due to the provider's deficient testing and training of the AI model (the second stage). However, should the provider and procurer fail to delineate their respective responsibilities at all, it is suggested that there should be a default rule (that can be opted out of) requiring the party that is best placed to assume the relevant responsibility. This is likely to be both provider and procurer at stage one; the provider and to a lesser extent the procurer at stage two; and the procurer at stage three.

---

[48] Panesar, *supra* note 33.

Finally, provided that B2B AI transactions are contracts for the sale of goods (or are even work and materials contracts[49]), the demarcation of spheres of responsibility between the provider and procurer should also make reference to matters such as whether the provider can sue for the price[50] and the incidence of risk of loss/damage[51] to the AI system, issues which may turn upon whether and when the property in the AI system passes.[52] The time at which property in the AI system passes depends on the contracting parties' intention, subject to the AI system first becoming existing if the contract is for future goods.[53] Consider the example of AI tax solution sold by Deloitte, which is likely to be a contract for future goods. This is because the construction of the AI system depends in part on the data to be selected and fed by the procurer (drawn from its payroll, accounts payable, and general ledger), as well as the testing and training of the AI model by the provider based on the data supplied by both the provider and procurer. Future goods come into existence when they are produced by the seller.[54] In the B2B AI context, the provider and procurer are not precluded from contractually agreeing, in a contract for the sale of an AI system to be constructed, that the property in the inchoate system will pass as it is added to from time to time.[55] But absent such contractual intention, there is a strong presumption against the passing of property in incomplete work,[56] and thus, the provider may not be able to sue for the price as the procurer might argue that property had not passed.[57]

However, assuming that B2B AI transactions are contracts for the supply of services, the above passing of property analysis will not apply. Instead, issues including whether the supplier of the AI service can sue for the price will depend in part on whether the contract is entire[58] or divisible. For example, in the case of Deloitte tax solutions, if the procurer engages Deloitte to supply an AI tax system for a certain amount of money on satisfactory completion of the whole system (or

---

49 It has been observed that while there is a shortage of authority on the passing of property from transferor to transferee in work and materials contract, the rules of the *SGA* are likely to apply: MG Bridge, *The Sale of Goods*, 4th ed (Oxford: Oxford University Press, 2019) at para 3.06.

50 *SGA*, s 49(1).

51 *SGA*, s 20.

52 Bridge, *supra* note 49 at para 3.01. But if the contract is structured as a license to use AI software, the passing of property analysis is unlikely to apply: see *infra* notes 62, 63 and accompanying text.

53 Bridge, *ibid* at para 3.29.

54 *SGA*, s 5(1).

55 By analogy with a contract for the sale of a ship to be built: see *eg*, *Clarke v Spence* (1836) 4 A & E 448; *Seath v Moore* (1886) 11 App Cas 350 (HL) [*Seath v Moore*]; *Reid v MacBeth & Gray* [1904] AC 223 (HL); *Sir James Laing & Sons Ltd v Barclay, Curle & Co Ltd* [1908] AC 35 (HL) [*Sir James Laing v Barclay*]. But the validity of this analogy depends on whether the AI systems discussed in this paper are goods, which will be analysed in Part IV(B). There is also the difficult issue, outside this paper's scope, of whether the data (which is fed into the AI system) in itself, is property, but the answer is probably not. The law distinguishes between the information itself (which includes data), the physical medium on which the information is recorded, and the rights to which the information gives rise; whilst the physical medium and the rights are treated as property, the information itself is not: *Your Response Ltd v DataTeam Business Media Ltd* [2014] EWCA Civ 281 at para 42 (Floyd LJ); *OBG Ltd v Allan* [2007] UKHL 21 at para 275 (Lord Walker).

56 *Seath v Moore*, *ibid*; *Sir James Laing v Barclay*, *ibid*.

57 *SGA*, s 49(1).

58 For an example of an entire contract, see *Cutter v Powell* (1795) 6 Term Rep 320 (KB).

part-payment for defined parts), the procurer need not pay, and Deloitte cannot sue for part-payment, if Deloitte fails to supply the whole system (or its stipulated part). Thus, the suggested solution is for Deloitte (or any provider) to contractually agree that part payment in the form of installments should be made at different stages of the AI system (using the three-stage framework as a guide). Or parties may consider adding a substantial performance[59] clause in the contract, stating that as long as the provider has substantially performed its obligations (with a clear specification of what amounts to substantial performance), the procurer is required to pay after deducting any loss suffered from incomplete or defective performance by the provider. In short, insofar as the procurer has agreed to pay for the services rendered, the provider can sue for the price, even if the end product (*ie*, the final AI system) is not delivered. But this is not the case in sale of goods, where the price is generally payable when property has passed to the buyer and upon delivery and acceptance.

The above analysis underscores the important issue of whether B2B AI transactions are contracts for the supply of services, sale of goods, or are work and materials contracts, as well as whether the AI systems examined in this paper are goods, all of which will be examined in the next section.

IV. Liability in the B2B Context

A. *Contract for Supply of Services or Sale of Goods, or Work and Materials Contract?*

Before the specific requirements under the statutes are examined, the question of whether the *SGSA* or *SGA* applies to B2B AI transactions, namely, whether these transactions are contracts for supply of services or for sale of goods (or even work and materials contracts), has to be addressed.

Based on the nature of the AI systems (examined in Part II) and the roles played by both the AI solutions provider and procurer (examined in Part III), a stronger case based on doctrinal and policy grounds can be made for the position that B2B AI transactions should be regarded as contracts for the supply of services, not for the sale of goods.

From a doctrinal perspective, the test is arguably whether the contract's substance is for the production of something to be sold by the seller to the buyer (a sale of goods under the *SGA*), or for the exercise of skill and labour to produce something wherein the provision of materials is ancillary (a supply of services under the *SGSA*).[60] It has been observed that the test is difficult to apply and the cases seem irreconcilable.[61] But, applying the test to the AI B2B context, the contract's substance seems to be for the skills and labour exercised by both the AI solutions provider and procurer in producing the AI system. In the Deloitte example, the AI

---

[59] See *eg*, *Hoenig v Isaacs* [1952] 2 All ER 176 (EWCA); *Dakin (H) & Co Ltd v Lee* [1916] 1 KB 566 (EWCA).

[60] *Robinson v Graves* [1935] 1 KB 579 (EWCA).

[61] Ewan McKendrick, *Goode on Commercial Law*, 5th ed (LexisNexis, 2016) at para 7.21.

solutions provider played a critical role in curating and feeding data into the AI system, and training and testing it afterwards. On the procurer's side, the AI system may depend upon the procurer feeding the system with its own proprietary data. Moreover, providing materials is ancillary: this paper examines AI systems essentially based on machine learning, whose learning routes are ultimately coded, and hence more akin to software.

Another doctrinal reason is that insofar as the agreement is structured as a licence to use AI software, there may be no transfer of property in goods, which the *SGA*[62] requires for establishing sale of goods contracts. This is because the agreement may contain terms and conditions vesting in the supplier the "title, copyright and all other proprietary rights in the software", or requiring the customer at the supplier's request to return the software or destroy parts of the software covered by the licence.[63] But this argument cannot be pressed too far: in the B2B AI transaction context, it would be difficult to vest proprietary rights in the AI system exclusively in the AI solutions provider. This is because AI solutions providers and procurers may be jointly responsible for data input and for AI system design and testing. If data rights (including access and use) are regarded as proprietary, the role of procurers in data management makes it impracticable for AI solutions providers to reserve all intellectual property rights. That said, where the licence does not vest proprietary rights in the supplier and the contract is structured as a perpetual licence granting the customer unfettered ability to use the items forever subject to copying restrictions and breach-related conditions, a first instance authority stated that the agreement is a sale of goods, albeit in the context of a different statute.[64]

From a policy perspective, it is more appropriate to subject AI solutions providers to the duty to exercise reasonable care and skill, *ie*, to apply the negligence standard (for the supply of services under the *SGSA*),[65] rather than the strict liability standard (for the sale of goods under the *SGA*). This is arguably appropriate as imposing strict liability on AI solutions providers omits consideration of the procurer's important role in the three-stage framework. Both procurers and AI solutions providers manage data, test and implement systems, and manage external threats.[66] Thus, it seems more proportionate as between the procurer and the provider that the standard of negligence be applied to what the provider is

---

62 *SGA*, s 2(1).

63 *Southware LBC v IBM UK Ltd* [2011] EWHC 549 (TCC) at para 95.

64 *Software Incubator Limited v Computer Associates UK Limited* [2016] EWHC 1587 (QB) at para 62. Regulation 2(1) of the *Commercial Agents (Council Directive) Regulations 1993* SI 1993/3053 applied here (*ie*, not the *SGA*). The Court of Appeal did not challenge the first instance judge's reasoning and conclusion that the perpetual license amounted to a sale of goods.

65 Bridge, *supra* note 49 at para 2.24. But note that if the contract is one of transfer of goods under the *SGSA*, the strict liability standard pertaining to the implied terms of quality and fitness applies: s 4(2) and s 5 *SGSA*.

66 NVIDIA, "Massachusetts General Hospital Use Artificial Intelligence to Advance Radiology, Pathology, Genomics" (5 April 2016), *NVIDIA Newsroom*, online: <https://nvidianews.nvidia.com/news/nvidia-massachusetts-general-hospital-use-artificial-intelligence-to-advance-radiology-pathology-genomics>.

responsible for.[67] This lets the procurer's contributory negligence be considered where relevant.[68]

However, the procurer as claimant would want to characterise the contract as a sale of goods (under the *SGA*) rather than a provision of a service (under the *SGSA*) because that characterisation places the AI solutions provider in a position of greater legal responsibility and potential liability.

Alternatively, the procurer may argue that B2B AI transactions are work and materials contracts that differ from pure services because the goods supplied are integral and not merely ancillary to the supplier's performance, and the labour and goods are inseparably blended.[69] This can amount to both a contract for services and a contract for the transfer of goods under the *SGSA*,[70] thereby triggering the strict liability standard that applies to implied terms of quality and fitness to which contracts for the transfer of goods are subject,[71] to the procurer's advantage. But this depends on whether the AI systems discussed in this paper are goods and it is argued below that on balance, this is unlikely to be the case.

B. *Is AI Software Goods?*

Do the AI systems deployed in taxation, stock trading, sales and marketing, and healthcare, which are software,[72] amount to goods? Such AI systems must be contrasted with hardware ones such as self-driving cars and robotics. Nothing in the *SGA*'s definition of goods[73] indicates whether AI systems are goods. The position of software under the *SGA* is a useful starting point. Computer hardware is indisputably goods;[74] so is hardware containing software.[75] The difficulty is whether standalone software, specifically the AI systems examined here, are goods. Although the case law has not arrived at a definitive position, the policy justifications suggest that software is not goods.

---

[67] But apportioning liability between sellers' strict liability under *SGA*, s 14(3) and buyers' negligence is not permitted: *Law Reform (Contributory Negligence) Act 1945* (UK), c 28.

[68] Arguably, buyers' contributory negligence can be considered as buyers cannot rely on the fitness term if they know of defects yet continue to use the goods: *Lambert v Lewis* [1982] AC 225 (HL) (but contributory negligence can break causation chain emanating from sellers' breaches: Michael G Bridge, "Defective Products, Contributory Negligence, Apportionment of Loss and the Distribution Chain" (1981) 6:2 Can Bus LJ 184). However, in AI B2B contexts, the issue is less about the awareness of defects and more about procurers' roles in AI system design and implementation.

[69] Bridge, *supra* note 49 at para 2.58. For an example of a mixed contract, see *Hyundai Heavy Industries Co Ltd v Papadopoulos* [1980] 1 WLR 1129 (HL).

[70] *SGSA*, s 1(1); Bridge, *ibid*.

[71] *SGSA*, ss 4(2), 5.

[72] See *eg*, IBM's Watson for Oncology (software letting doctors quickly retrieve patients' information, and speedily access different medical treatments and their pros and cons tailored to patients' conditions)—*IBM*, "Watson Health", online: <https://www.ibm.com/products/clinical-decision-support-oncology>.

[73] *SGA*, s 61(1).

[74] *Rubicon Computer Systems Ltd v United Paints Ltd* [2000] 2 TCLR 453 (EWCA).

[75] *St Albans City and District Council v International Computers Ltd* [1996] 4 All ER 481 (EWCA) [*St Albans City*]; *Southwark LBC v IBM UK Ltd* [2011] EWHC 549 (TCC); *Toby Constructions Products Pty Ltd v Computer Bar (Sales) Pty Ltd* [1983] 2 NSWLR 48 (NSWSC).

From a doctrinal perspective, there are cases that support the proposition that software is not goods. For example, it has been held that standalone software is not goods, unless the software is embodied within or supplied as hardware.[76] Consistent with this reasoning, albeit in the context of a different legislation,[77] the Court of Appeal, having surveyed the cases in the UK, other common law jurisdictions and Europe, held that because goods involve tangible property, computer software supplied by electronic means and not on any tangible medium is not goods as it does not involve tangible property.[78] The distinction between the supply of software without any tangible media (which is not goods) and with such medium (and is thus goods) is also supported by other cases.[79] However, the Attorney-General of the Court of Justice of the EU held that electronically supplied software is goods under the legislation in light of its wording, context and objectives, which do not restrict the meaning of goods only to tangible items, and can also apply to intangible ones.[80]

However, there are cases (not discussed by the Court of Appeal) that seem to suggest that software is goods. For instance, one case assumed that the supply of computer software amounted to both a supply of goods and supply of services under the *SGSA*.[81] This implies that software is goods under the *SGA* because the *SGSA*[82] and *SGA*[83] define goods identically. Further, in one case, the court assumed the provision of software was a sale of goods under the *SGA*,[84] and in another case, assumed the software transaction was both a sale of goods under the *SGA* and a supply of goods under the *SGSA*.[85] Because the courts in these cases merely accepted the parties' (undisputed) submission that software is goods, it is not clear how much weight should be placed on these decisions.

From a policy perspective, a stronger argument can be made that software is not goods. Parliament created a separate protected category of digital content (including software) in the *Consumer Rights Act 2015*[86] ("*CRA*"), without similarly amending the *SGA* to include this category. This implies Parliament did not intend to change the status quo, namely, that digital content including software is not goods under the *SGA*. The *CRA* draws no distinction between software or digital content supplied in

---

76 *St Albans City*, *ibid*; *Gammasonics Institute for Medical Research Pty Ltd v Conrad Medical Systems* [2010] NSWSC 267.

77 *Commercial Agents (Council Directive) Regulations 1993*, *supra* note 64, reg 2(1).

78 *Computer Associates UK Ltd v The Software Incubator Ltd* [2018] EWCA Civ 518 at paras 25–40 [*Computer Associates*] (Gloster LJ). Note that the Supreme Court had referred two questions as to the meaning of the Directive to the Court of Justice of the EU: (1) whether electronically supplied software amounts to goods; and (2) whether the supply of software by way of a perpetual and/or limited term licence to use it amounts to the sale of goods.

79 See *eg*, *Beta Computers (Europe) Ltd v Adobe Systems (Europe) Ltd* [1996] SLT 604, at 608L–609B (CSOH) [*Beta Computers*] (cited with approval in *Computer Associates*, *ibid*).

80 Advocate General Tanchev, "Case C-410/19 The Software Incubator Ltd v Computer Associates UK Ltd", Opinion, (17 December 2020) at paras 51–78. The CJEU's ruling is pending, as is the Supreme Court's final decision.

81 *SAM Business Systems Ltd v Hedley and Co* [2002] EWHC 2733 (TCC); [2003] 1 All ER (Comm) 465.

82 *SGSA*, s 18.

83 *SGA*, s 61(1).

84 *Astea (UK) Ltd v Time Group Ltd* [2003] EWHC 725 (TCC).

85 *Kingsway Hall Hotel Ltd v Red Sky IT (Hounslow) Ltd* [2010] EWHC 965 (TCC).

86 (UK), c 15.

disks, and software or digital content directly downloaded.[87] In both cases, consumers are entitled to the same protection. But the *SGA* has not been correspondingly amended. One possible reason is the need to protect consumers who are end-users of software services. But in the B2B context, supply of software may be subject to testing before delivery, therefore possibly involving prior scrutiny by the procurer.

Another policy reason that may cast doubt on AI software being classified as goods is that where computer software is designed to perform tasks that might have once been performed by a professional supplying professional services, to whom negligence standard for liability applies, it is questionable that liability for non-performing software should be strict.[88] Insofar as AI software such as those related to taxation, healthcare and possibly stock trading perform the functions of professional service providers, it is suggested that liability should be fault-based.

C. *Why is the* SGA *Still Relevant?*

In sum, a stronger case can be made that B2B AI transactions are contracts for the supply of services, and the policy reasons indicate that AI software is unlikely to be goods, thus making the *SGSA* the applicable legislation and the standard of liability fault-based. However, the *SGA* (or its implied terms such as satisfactory quality or fitness for purpose which is subject to strict liability and which also apply to contracts for the transfer of goods under the *SGSA*[89]) remains potentially relevant for two reasons. First, in a contract for the supply of software, it has been held to be *sui generis*, thus leaving open the question of whether liability is strict or fault-based.[90] The more important reason is that while the strict liability standard under the *SGA* such as the implied requirements of satisfactory quality and fitness for purpose places most of the burdens on AI solutions providers (which is anomalous as the three-stage framework shows an interdependent buyer-seller relationship), judicial interpretation of these implied requirements in light of the framework will ameliorate the harshness of strict liability (as the next section will show). Accordingly, analysing how the framework should apply to the *SGA* (specifically the implied terms of quality and fitness, which also apply to the transfer of goods under the *SGSA*) remains relevant. The specific requirements in the *SGSA* and the *SGA* are examined below.

---

87 *CRA*, s 2(9).

88 Bridge, *supra* note 49 at para 2.25.

89 *SGSA*, ss 4(2), 5.

90 *Beta Computers*, *supra* note 79 at 396. While neither the *SGA* nor the *SGSA* was discussed in this Scottish decision, this case may be relevant to a judge who wishes to analogise AI software as *sui generis*. In this case, the software was not supplied in the form of a physical medium (such as a disk). But had it been done so, the software could be construed as goods; nor was the software classified as standalone, which would not be construed as goods (see *St Albans City*, *supra* note 75). Rather, the Court of Session said that "the subject of the contract was a complex product comprising the medium and the manifestation within it or on it of the intellectual property of the author… in the form of the program material contained. There are perhaps no true analogies of this type of product."

D. *Supply of AI Systems as Services, and Liability under the* SGSA

1. *Reasonable care and skill*

Under s 13 of the *SGSA*, where suppliers act in the course of a business, there is an implied term that suppliers carry out the service with reasonable care and skill, *ie*, "the [service] would be provided without negligence."[91] However, commentaries on this section avoid expositions of the law because precedents are likely less helpful in addressing the reasonable care and skill issue (as it will be largely resolved in the particular circumstances of each case).[92]

However, one authority provides an arguably helpful analogy on how s 13 can apply to the B2B AI context. In *Trebor Bassett Holdings Ltd v ADT Fire & Security Plc*,[93] the claimant sued the defendant for breach of contract and for breaching s 13 following a fire in its factory. The claimant made popcorn and other confectionaries; the defendant supplied a fire suppression system. In the popcorn production area of the claimant's factory, corn was popped in pans of oil before being transported via an elevator to a hopper for packaging. The defendant designed and installed an automatic "CO2 suppression" system in the elevator and the hopper. A fire broke out in the hopper. It spread, destroying the whole factory. The court found that the defendant was obliged to take reasonable skill and care in designing the system and that the defendant breached s 13.

The court made three key points. First, in designing the relevant system, the defendant must understand all relevant risks arising from the system's design and use in relation to its context. Second, it is important to identify the specific kind of harm (in this case the type of fire) the system supplied was meant to avert, and whether the system fulfilled the function. Finally, in deciding whether the standard of reasonable care and skill was met, the court reviewed the relevant context and specific factors (including the type and location of sensor detecting the fire).

The first point concerning the nature, scope and extent of risks has important implications for AI systems. As examined earlier, risks relating to the AI system can be distinguished based on which part of the three-stage framework they fall within. Risks can relate to data management including inaccurate, inadequate or non-representative data. Risks pertaining to systems development and implementation can

---

91 *Abramova v Oxford Institute of Legal Practice* [2011] EWHC 613 (QB) at para 58.

92 *Eg*, Richard Christou, *Sale and Supply of Goods and Services*, 3d ed (London: Sweet & Maxwell, 2015) at para 1.4.2.1.

93 [2012] EWCA Civ 1158 [*Trebor Bassett*]. The court concluded at paras 45, 48 that a contract for the installation of a fire suppression system was not to be equated with a contract for the supply of goods. Thus, the strict liability standard that applies to the implied terms of quality and fitness did not apply. It said that (a) while the system was made of various parts, what made it a system was the design, the selection and integration of all relevant information into a designated system to achieve a particular effect; (b) what the supplier agreed to supply was primarily design skills and care in exercising them; and (c) the shortcomings in the system were matters of design, not the inherent quality of the goods which were also supplied. By analogy, AI solutions providers arguably provide "design" service (*ie*, data management (first stage) and system development (second stage)) by integrating all relevant "information" (*ie*, how and which artificial neural networks and data were selected and inputted) into a specific system in order to attain a particular goal (such as increased efficiency and accuracy in tax compliance, increased sales and marketing, or improved detection and diagnosis of tumours).

materialise due to insufficient or improper training and testing of the AI system and usage error. Risks related to managing external threats may also be shared between the procurer and AI solutions provider (although the procurer is usually responsible). As both parties may share risks fundamentally affecting the AI system's ultimate performance, it will be important to delineate their respective or dominant responsibilities for these risks to see who may be negligent.

On the second point about the specific kind of harm, supplying the AI system may be to specifically address certain problems or achieve certain goals within particular contexts, for example the IBM Watson system's goal of detecting cancer tumours early and accurately for diagnosis and treatment. However, Watson has not performed properly, generating false positives and false negatives. Applying s 13 is complicated by procurers' and AI solutions providers' shared responsibilities in achieving AI systems' goals. Using the framework to ascertain procurers' and providers' relative responsibilities can help apportion liability.

On the court's third and final point, it found the defendant had insufficiently considered which sensor was most appropriate for fire detection and in which location in the hopper the sensor should be best placed, resulting in the fire going undetected. By analogy to AI systems, it must be ascertained whether procurers or providers made particular decisions affecting AI systems' performance. For example, in the framework's second stage (*ie*, system development and implementation), perhaps the optimal activation function[94] in an AI system using artificial neural networks was insufficiently considered. Activation functions are mathematical equations attached to each neuron, determining AI system's output using artificial neural networks. The AI solutions provider's wrong choice of activation function can engender wrong predictions by the AI system.

2. *Causation*

However, AI solutions providers may seek to evade liability by arguing that there is a lack of causation where the AI system's unexpected performance can be attributed to inherent unpredictability in machine learning routes.[95]

Three responses are warranted. First, the impossibility of establishing causation is overstated. Legally, the procurer only has to prove on the balance of probability (and not beyond a reasonable doubt) that it has a real chance of obtaining the benefit had the AI solutions provider not been negligent.[96] In other words, courts do not require procurers to prove on the balance of probability that but for the negligence of AI solutions providers, procurers would have actually obtained the benefit (as

---

[94] Dinesh Kumawat, "7 Types of Activation Functions in Neural Network" (6 Dec 2021), *Analytic Steps*, online: <https://www.analyticssteps.com/blogs/7-types-activation-functions-neural-network>.

[95] *Eg*, European Commission, Expert Group on Liability and New Technologies—New Technologies Formation, *Liability for Artificial Intelligence and Other Emerging Digital Technologies* (Brussels: EC, 2019) at 20–22 [EC, *Liability for AI*].

[96] The but-for test is applied to the issue of causation of some economic loss and then the loss must be ascertained in terms of value by considering the role of chance. The damages owing can be reduced by reference to contributory negligence/apportionment. Note also that loss of a chance is not the exclusive perspective in relation to possible causal inquiries.

opposed to the chances of obtaining it).[97] For example, in the sales and marketing context, if the AI system wrongly analyses customers' preferences resulting in lost sales, the court compensates procurers for the loss of that chance of obtaining the benefit. It should be easier for claimants to prove loss of chance of obtaining something rather than proving it would have obtained that thing itself.

Next, procurers' sustained losses could be attributable to a few causes, such as AI solutions providers' negligence, or machine learning's unexplainable workings (including the weight assigned by neurons in artificial neural networks), or even procurers' contributory negligence. The question would be to what extent providers' negligence reduced procurers' chances of obtaining the benefit. It may not be possible to assign precise probabilities to providers' negligence alongside the probability of highly complex and opaque artificial neural networks leading to sub-optimal performance. However, the lack of precise assignments of probability will not completely eliminate the role played by a provider's negligence (such as by feeding erroneous data or failing to correct mislabelled data) in reducing a procurer's chances. Thus, it is reasonable to take the position that a provider's negligence reduces a procurer's chances, to the extent determined by courts on a case-by-case basis bearing in mind the parties' contractual intentions, market practice, and all relevant circumstances.

Second, taking the black box nature of AI systems to its logical conclusion, no AI systems supplier could ever be held liable because causation could never be established, effectively insulating suppliers from liability. Further, accepting the impossibility or unlikelihood of demonstrating causation in machine learning undermines efforts[98] to increase AI system suppliers' (and programmers') accountability because liability—a key mechanism for vindicating rights and deterring wrongdoing—would be ineffectual.

Finally, where AI solutions providers fall below their standard of care in any way depending on the three-stage framework, the unpredictabilities in machine learning are unlikely to break the chain of causation. This is because for a chain of causation to be broken, the law requires a positive act (committed by either claimants or third parties) that is deliberate (*ie*, intentional), voluntary (*ie*, not done under duress or lack of mental capacity), informed (*ie*, done while understanding the consequences) and unreasonable.[99] For example, a defendant AI solutions provider would find it difficult to argue that a claimant procurer's failure to extract and input accurate and complete corporate data into the AI system (insofar as such data are needed from the procurer) broke the chain of causation. This is because omission, even if wrongful, deliberate, informed and unreasonable, will not break the chain of causation.[100] But an AI solution provider can argue that at stage one of the framework, a procurer's deliberate and unreasonable inputting of inaccurate data into the AI system would break the chain of causation leading to the procurer's own losses, as the procurer

---

97 *Brown v Ministry of Defence* [2006] EWCA Civ 546 [*Brown*].

98 *Eg*, Ronan Hamon, Henrik Junklewitz & Ignacio Sanchez, *Robustness and Explainability of Artificial Intelligence* (Luxembourg: Office of the European Union, 2020).

99 Nicholas J McBride & Roderick Bagshaw, *Tort Law*, 6th ed (London: Pearson, 2018) at 297, 298.

100 *Ibid* at 299.

would have been chiefly responsible for bringing about sub-optimal performance of the AI system.

3. *Contributory negligence*

In addition to helping determine whether a defendant provider is negligent, the framework can also assist in determining whether the claimant procurer is contributorily negligent.[101] In *Trebor Bassett*, although the court found that the defendant failed to exercise reasonable care and skill in designing the fire suppression system, it agreed with the trial judge that the claimant was contributorily liable for 75% of the losses, having failed to segregate the production area (where the fire broke out) from the rest of the building and failing to install sprinklers.[102] Analogising to the AI systems transaction context, even where AI solutions providers fail to ensure accurate data are selected and fed into the AI system (stage one), or they have been negligent in testing at the system development stage (stage two), procurers can be contributorily liable by feeding inaccurate or incomplete corporate data into the system at the data management stage. Or at the system implementation stage, the procurer may have failed to question or override AI system decisions despite, for example, their deviation from well-established views or internal protocols. To be contributorily negligent, procurers must have contributed to their own loss without breaking the chain of causation. Because B2B transactions concern economic losses, and the parties have joint responsibility under the three-stage framework, courts should adopt the loss of chance approach to causation discussed above.[103]

4. *Suggested solutions*

It was recommended in Part III that procurers and AI solutions providers should maintain joint data movement and joint data description records. These enable the allocation of responsibilities as they show the identification and nature of risks in relation to data management, system implementation and development, and external threat management, and indicate which party took what steps to mitigate them. This proposal is similar to the mandatory "log" proposal for ascertaining who took which decisions, and whose negligence can be regarded as causative or principally causative, in accidents caused by autonomous vehicles.[104] In sum, the three-stage framework enables the application of statutory negligence law to the AI context if AI systems are regarded as services, thereby allowing the allocation of responsibilities and liabilities to be mapped out.

Having analysed the liability issues related to the *SGSA* in light of the framework, the next section examines the *SGA*.

---

[101] *Law Reform (Contributory Negligence) Act 1945*, *supra* note 67.
[102] *Supra* note 93.
[103] *Eg*, *Brown*, *supra* note 97.
[104] EC, *Liability for AI*, *supra* note 95.

E. *Strict Liability under the* SGA *for AI Systems as Sale of Goods (or under the* SGSA *for AI Systems as Transfer of Goods)*

1. *Satisfactory quality*

Section 14(2A) of the *SGA* defines goods to be of satisfactory quality "if they meet the standard that a reasonable person would regard as satisfactory, taking into account of any description of the goods, the price (if relevant) and all the other relevant circumstances". Section 14(2B) lists non-exhaustive factors in assessing goods' quality: state and condition, fitness for all purposes for which the goods of the kind are commonly supplied, appearance and finish, freedom from minor defects, safety and durability.

What amounts to satisfactory quality in the context of AI systems cannot be answered in the abstract; it has to be analysed in the context of the specific contract.[105] It has been observed that "… satisfactory quality is a subtle and varying standard whose application is clearly influenced by the type of damage caused by defective goods, as well as the nature of the goods."[106]

Because what amounts to satisfactory quality has to be also determined by all the relevant circumstances, it is important to consider what this means in the AI context. To illustrate, consider an AI tax solution conceptualised by Deloitte. The AI system's objective is to assist the company's compliance with a specific employment tax obligation more accurately, quickly, and cost-effectively. Deloitte first feeds data including all relevant rules and regulations, training materials and dictionary into the AI system. Deloitte then works with the company (*ie*, the procurer) to select, analyse and feed the company's data (drawn from the payroll, accounts payable, and general ledger) into the system. Next, in training and monitoring the AI system, the software dashboard displays curated and labelled data. Deloitte and the company (both AI solutions provider and procurer) correct missing, inaccurate, or mislabelled data as appropriate. After system development, the procurer can deploy the software. Given the importance of data, it is suggested that the relevant circumstances should include whether AI solutions providers keep proper records of data movement that identify which data were fed into the AI system and from where as well as data description records of the content and quality of datasets. The absence of, or material deficiencies in, such records might give rise to the inference that the AI system is not of satisfactory quality.

However, if the AI solutions provider drew the procurer's attention to any matter rendering goods unsatisfactory before making the contract, the provider need not comply with the satisfactory quality requirement.[107] The provider need not highlight

---

[105] Bridge, *supra* note 49 at para 7.49.

[106] *Ibid* at para 7.48.

[107] *SGA*, s 14(2C)(a). Note that s 14(2C)(b), which is concerned with the buyer examining the goods before the contract is made, is unlikely to apply to the context of B2B AI transactions because the provider would not have the goods, *ie* the AI system, completed and ready, before the contract was entered into between the provider and procurer. After all, the provider needs the cooperation of the procurer in selecting and feeding its own data into the AI model in order to construct the AI system, as shown in the Deloitte AI tax solution. But if there is a more or less ready-made AI system (without the need for the procurer's data), such as the supercomputer in the case of *Tyndaris*, *supra* note 7, the exception will not

the exact defect to the procurer because the section only requires the seller to draw the buyer's attention to matters rendering goods unsatisfactory.[108] Thus, if Deloitte explicitly informed the company that inaccurate, incomplete or non-representative data (the first stage of the framework) will make the AI system's quality unsatisfactory, regardless of whether the data are analysed or fed by Deloitte or the procurer, the satisfactory quality requirement will arguably not apply to Deloitte. Or Deloitte may provide that its testing and training responsibilities (the second stage) do not extend beyond certain parameters. How should courts construe such qualifications? Specifically pointing out goods' physical abnormalities will not excuse sellers from the satisfactory quality requirement if the abnormalities in themselves do not show the goods' latent defects.[109] Therefore, it seems that the court will absolve sellers only where buyers accept the qualifications with truly open eyes. Analogously, drawing the procurer's attention to the black box nature of machine learning in AI software should not exempt the AI solutions provider from the satisfactory quality requirement. The provider's qualifications must be much more precise. Because procurer and provider are likely involved in intense relational dynamics in developing AI systems, the contract should precisely delineate the provider's responsibility for satisfactory quality.

2. *Fit for purpose*

Turning to the strict liability requirement to sell goods reasonably fit for the purpose the procurer expressly or impliedly informs the provider of,[110] the disclosure of purpose by the buyer to the seller is not in law restricted to special or esoteric purposes. It extends to common purposes.[111] Thus, insofar as the purpose is commonplace, the buyer need not expressly communicate it to the seller.[112] Applying the law to the Deloitte example above, the corporate procurer need not inform Deloitte that its purpose of buying the software is to comply with tax obligations by ensuring all submissions to the tax authority are accurate and complete. The seller is entitled to assume that that is the procurer's purpose. But note that because the implied term of fitness of purpose does not refer to absolute fitness, but "reasonable" fitness, an AI system may satisfy this term even if it does not quite achieve the outcome expected by the procurer in light of the purpose(s) made known.

Moreover, the seller will not be liable if the buyer does not rely, or if it is unreasonable for it to rely, on the seller's skill or judgment.[113] But if the seller was given an opportunity to exercise skill or judgment, the buyer would be entitled to protection from not only defects that the seller should have detected, but also "defects …

---

apply if the procurer did not examine the goods, and even if the procurer had done so, will apply only if the defect ought to have been revealed by the examination: *Macdonald v Pollock* [2011] CSIH 12, 2013 SC 22 at para 34.

[108] Bridge, *supra* note 49 at para 7.85.

[109] *Stephenson v Cookson* [2009] EWCA Civ 1270.

[110] *SGA*, s 14(3).

[111] *Ibid*.

[112] *Preist v Last* [1903] 2 KB 148; *Hamilton v Papakura District Council* [2002] UKPC 9.

[113] *Ibid*.

latent in the sense that even the utmost skill or judgment on the part of the seller would not have detected them."[114] In the context of B2B AI sales, both AI solutions provider and the procurer may exercise skill and judgment in the data management and system development stages. Thus, the *SGA*'s assumption regarding parties' relative inequality may not hold. However, Deloitte, the AI solutions provider, is arguably ordinarily regarded as the tax consultant and expert. Hence the procurer should be able to rely on Deloitte inputting correct and updated data (including tax rules and regulations) into the AI system. The procurer is however the likely expert in managing its data categories for different employee types, so it would be unreasonable to rely on Deloitte for data accuracy that potentially affects the performance of the whole system. Thus, one key question is when and to what degree procurers may reasonably rely on AI solutions providers' skill and judgement.[115] This depends on procurers' and providers' precise roles in each of the three stages in the proposed framework, which vary with different AI systems.

To conclude, despite the strict liability requirements in the *SGA*, interpreting the law in light of the framework provides a nuanced and fairer approach towards calibrating the AI solutions provider's potential responsibility and liability. But parties have to be mindful of exactly how they delineate and divide their responsibilities in AI design and development (particularly the first and second stages of the framework) in their agreements.

Parties may, however, behave defensively in order to shift risk to the other. In the majority of B2B transactions where the parties have generally equal bargaining power and are represented by lawyers, agreements are likely to contain an entire agreement clause (which states that the agreement contains the totality of the parties' bargain) and non-reliance clause (which states that the parties have not relied on any representations, statements or claims in entering into the contract), the result of which will be that the seller's (AI solutions provider's) liability is precluded or restricted. So the effect of exclusion clauses must be examined.

F. *The Impact of Exclusion or Limitation of Liability Clauses*

1. *Exemption/limitation clauses*

AI solution providers may resort to three techniques for excluding or limiting their liability. First, they may ensure their representations or claims are not part of the contract by including a broadly drafted entire agreement clause[116] (but the latter will not preclude courts from implying terms into the contracts).[117] Second, they can include a broadly drafted non-reliance clause.[118] Finally, they can include exemption clauses in the contract expressly excluding themselves from liability.

---

[114] *Henry Kendall & Sons v William Lillico & Sons Ltd* [1969] 2 AC 31 at 84 (HL).

[115] Partial buyer reliance suffices: *Cammell Laird & Co Ltd v Manganese Bronze and Brass Co Ltd* [1934] AC 402 at 427 (HL).

[116] *NF Football Investments Ltd v NFCC Group Holdings* [2018] EWHC 1346 (Ch).

[117] *J N Hipwell & Son v Szurek* [2018] EWCA Civ 674.

[118] *Springwell Navigation Corporation v JP Morgan* [2010] EWCA Civ 1221.

For example, AI solution providers can expressly limit or exclude liability arising from incomplete, inaccurate or unrepresentative data; likewise they can also limit or exclude liability arising from inadequate or flawed AI system training and testing. The exemption clause can exclude liability not only for damages to procurers, but also consequential[119] losses arising from AI system use, operation or maintenance.

The second technique concerning non-reliance clauses, insofar as they exclude liability for misrepresentation rather than define parties' obligations, will be subject to the reasonableness requirement under s 11 of the *UCTA*, as stated under s 3 of the Misrepresentation Act 1967.[120] The third technique will also be subject to the same *UCTA* requirement.[121] It is trite law that excluding liability resulting from the breach of the *SGA* implied terms of satisfactory quality and fitness for purpose is subject to the *UCTA* reasonableness requirement;[122] so is excluding liability resulting from the breach of the *SGSA* reasonable care and skill requirement.[123] In determining whether the exclusion is reasonable, courts consider various guidelines including (a) parties' bargaining power; (b) whether the buyer was induced to accept the term; (c) whether the buyer knew or ought reasonably to have known of the existence and extent of the term; (d) whether the conditions imposed on making the claims are reasonable;[124] (e) whether the goods were manufactured, processed or adapted to the customer's special orders; and (f) the availability, affordability and efficiency of insurance.[125] The question is how the guidelines should be analysed in the B2B AI transaction context. It is suggested that courts should apply the guidelines with reference to the proposed three-stage framework. Overall, as shown below, the exemption clauses could be construed as unreasonable under the guidelines.

Guideline (a) likely advantages sellers (*ie*, AI solutions providers) in B2B transactions, and will find favour with judges supporting freedom of contract and party autonomy. This may mean upholding the exclusion clause where there is equality of bargaining power between commercial parties,[126] as is arguably the case in B2B transactions. But this is not decisive for two reasons. First, courts have struck down non-reliance clauses for violating the reasonableness requirement even where parties have equal bargaining power and legal representation, for example on the basis that the defendant would have incurred liability for misrepresentations but for the non-reliance clause.[127] The courts have mentioned the importance of pre-contractual inquiries (*ie*, representations made in conveyancing) but nothing in the judgments suggests such reasoning will be restricted to that context.[128] Arguably, pre-contractual inquiries gain even greater importance in the AI context given AI

---

[119] *Eg*, *Hotel Services Ltd v Hilton International Hotels (UK) Ltd* [2000] BLR 235 (EWCA).

[120] *First Tower Trustees Ltd v CDS (Superstores International) Limited* [2018] EWCA Civ 1396.

[121] *UCTA*, s 3.

[122] *UCTA*, s 6(1A)(a).

[123] *SGSA*, s 2(2).

[124] *UCTA*, Schedule 2.

[125] *UCTA*, s 11(4).

[126] See *eg*, *Watford Electronics Ltd v Sanderson CFL Ltd* [2001] 1 All ER (Comm) 696 at para 55 (EWCA); *Goodlife Foods Ltd v Hall Fire Protection Ltd* [2018] EWCA Civ 1371 at para 103. *Cf Britvic Soft Drinks Ltd v Messer UK Ltd* [2002] 1 Lloyd's Rep 20 at 57, 58 (EWCA).

[127] *First Tower Trustees Ltd v CDS (Superstores International) Limited* [2018] EWCA Civ 1396.

[128] *Ibid*.

systems' complexity. Second, courts can examine procurers' and AI solutions providers' joint data movement and data description records to identify who has the expertise, in order to determine whether the provider can reasonably exclude liability. Unequal technological or specialist expertise between a procurer and a provider may render a provider's exclusion of liability arising from a machine learning system unreasonable.

Guideline (c) may favour procurers. Given AI systems' highly technical nature, unless AI solutions providers specifically draw procurers' attention to exemption clauses, procurers may not know of them and their implications. Even if the procurer is sophisticated in the sense of knowing the clauses and has been involved in the data management aspect of the three-stage framework, the procurer would not necessarily fully understand the limits of the provider's responsibilities in the complex relational dynamics of their shared responsibilities. For example, consider an AI taxation system. While the procurer may be involved in the data management part of the framework insofar as corporate data have to be extracted from the company's general ledger and payroll, it may not be involved in training and testing the AI system. In this manner, it will be unreasonable for the provider to exclude liability in blanket terms including data management, and system development and implementation.

On guideline (d), because AI systems are complex (such as one using machine learning, particularly artificial neural networks), an exclusion term will likely be construed as unreasonable if the AI solutions provider imposes a short timeframe on the procurer to make claims.

On guideline (e), while the goods or services being adapted to the buyer's special order may render sellers' exclusions of liability reasonable because this may entail special risks, courts should be cautious in doing so in the AI systems transaction context. Even if customisation entails particular risks, the complexity of AI systems may necessitate both procurers' and AI solutions providers' involvement in the three-stages of the framework, and customisation does not mean that the procurer is in a better risk-bearing position. It would be unreasonable to accept that the procurer knows or accepts all risks associated with AI system performance.

On guideline (f), the insurance factor may be significant. The reasonableness of the exemption clause depends on whether the insurance is potentially available to either party at the time of the conclusion of the contract and its relative cost, and not the actual insurance taken by the parties.[129] There is much academic debate on whether using AI systems such as autonomous vehicles should be accompanied by mandatory insurance cover, particularly in relation to third-party liability.[130] Analogously, it remains to be seen how the insurance market will respond to the demands of AI solutions providers and procurers concerning the extent of the insurance coverage and pricing of premiums, after taking into account which party is better placed to address the risks and the legal validity of liability exclusion clauses.

---

[129] *The Flamar Pride* [1990] 1 Lloyd's Rep 434 (QB).

[130] Andrea Bertolini, "Insurance and Risk Management for Robotic Devices: Identifying the Problems" (2016) 16:3 Global Jurist 291.

In sum, on balance, a stronger argument can be made that exclusion clauses are unlikely to be held reasonable based on an application of the guidelines. But it does not follow that AI solutions provider cannot limit liability; on the contrary, a carefully drafted limitation clause is more likely to pass the reasonableness test than a blanket exclusion clause such as one that exempts the AI solutions provider from liability in all three stages in the three-stage framework. Another advantage of a carefully drafted limitation clause is that it avoids the problem of courts generally having no power to sever unreasonable parts of the clause from the reasonable parts.[131] So, for example, a clause that states that the provider's liability will be restricted to a specific sum of money[132] that rests on evidence-based reasons, or that the provider will be held liable neither for the data that is selected and fed into the AI system by the procurer nor for improper operation of the AI system, is more likely to pass the reasonableness test than one that states that the provider will not be held liable for any data that is fed into, or for any testing of, the AI system.

2. *Duty-defining clauses*

The above analysis assumes that the clauses exclude or restrict liability. However, AI solutions providers can draft clauses defining parties' obligations, rather than exempting themselves from liability.[133] For example, a provider can undertake a duty to deliver an AI system produced with reasonable care and skill, or one of satisfactory quality and fit for purpose regarding the data selected and fed into it by the provider itself, but specifically exclude any duty to prevent or detect defects or malfunctions connected to the testing, training and deployment of AI systems.

On the one hand, such duty-defining clauses are arguably disguised exemption clauses caught by s 13(1) of the *UCTA* which prohibits excluding or restricting liability by "reference to terms and notices which exclude or restrict the relevant obligation or duty." On the other hand, such a broad interpretation of s 13(1) seriously threatens freedom of contract in B2B transactions where parties generally have equal bargaining power.

One way of addressing this issue is to apply the but-for test: whether liability will arise but for the challenged clause.[134] It has been held in negligence cases that as long as the clause does not shift liability to the buyer (*ie*, the victim of negligence), then it is a duty-defining clause.[135] But if it does, it will be construed as an exemption clause.[136] In the above example, because the purported duty-defining clause shifts risk to procurers for any defects of the AI system in connection

---

[131] *Stewart Gill Ltd v Horatio Myer & Co Ltd* [1992] QB 600 but note that where the different parts are independent of each other, courts may sever (*Regus (UK) Ltd v Epcot Solutions Ltd* [2008] EWCA Civ 361).

[132] But if there is inequality of bargaining power, the court is likely to hold that the clause is unreasonable: *St Albans City*, *supra* note 75.

[133] *Springwell Navigation Corporation v JP Morgan* [2010] EWCA Civ 1221.

[134] *Smith v Eric S Bush (A Firm)* [1990] 1 AC 831 (HL).

[135] *Thompson v T Lohan (Plant Hire) and Hurdiss (JW)* [1987] 1 WLR 649 (EWCA).

[136] *Phillips Products v Hyland* [1987] 1 WLR 659 (EWCA).

with its testing, training and deployment, it will be construed as an exemption clause and will be subject to the reasonableness requirement. However, a broad interpretation of the "but for" test renders all duty-defining clauses exemption clauses. It has been suggested that s 64(1) of the *CRA* may offer guidance: a term in a consumer contract may not be assessed for fairness insofar as it specifies the contract's main subject matter or the assessment relates to the contract price's appropriateness,[137] the rationale being that the consumer likely knows, negotiated and agreed to the main subject matter. Thus, only non-core terms will be subject to fairness review.

But the question of what is meant by the contract's "main subject matter" remains. In cases relating to supplying an AI system for sales and marketing, taxation, healthcare and share trading, one view is that the contract's "main subject matter" is the AI system itself and the specific functions that it is supposed to perform or the specific goals that it is supposed to attain. Implicit in the main subject matter—supply of the AI system—includes data management as well as system development and implementation (the first and second stages of the framework). But given the complexity (and black box nature) of the AI system, it does not seem reasonable to assume procurers have accepted all aspects of the three-stage framework as "main subject matter". This will be too over-encompassing. Hence, the "main subject matter" test should not apply to AI systems. But to argue that aspects of the three-stage framework such as data management and system development escape the contract's "main subject matter"—AI systems—seems artificial too. Nevertheless, even if a broad interpretation of the "but for" test is adopted, so that a duty-defining clause is construed as exempting liabilities, the defendant is not unfairly disadvantaged: courts must still subject the clause to the reasonableness requirement under the *UCTA*.

V. Conclusion

As the market for the design and supply of AI systems is growing rapidly, B2B disputes and private liability relating to AI system transactions will need attention. This article focuses on the allocation of responsibility between AI solutions providers and procurers, disavowing the assumption that AI solutions providers will necessarily be the only ones to face liability, or that no party will be liable due to the 'autonomous' nature of machine learning systems.

B2B AI transactions raise challenging private law issues because of the nature of AI systems and the new relational dynamics between AI systems providers and procurers. Thus, this article advances a three-stage AI design and development framework, comprising data management, system development and implementation, and external threat management. It is argued that the three-stage framework should be integrated into an analysis of liability in commercial law, in order to provide a clearer understanding of the legal risks and positions.

---

[137] Mindy Chen-Wishart, *Contract Law*, 6th ed (Oxford: Oxford University Press, 2018) at 433.

Finally, the analysis in this article reflects private law's in-built adaptability to technological developments. Reforming the law to cope with AI is not necessarily a solution for the challenging issues raised by new technologies, although there is scope for such consideration. The need to consider new normative doctrines such as the recognition of autonomy in AI systems is not ruled out,[138] but this is a matter for another work.

[138] *Eg*, Simon Chesterman, "Artificial Intelligence and the Limits of Legal Personality" (2020) 69:4 ICLQ 819; Nadia Banteka, "Artificially Intelligent Persons" (2020) 58 Houston L Rev.