

# DREWTECH SERIES CHAPTER 4 18 February 2020

## DIABOLUS EX MACHINA - Artificial (Un)intelligence and Liability

### INTRODUCTION

A self-driving car fails to “see” a truck blocking the road, and collides with it at high speeds, shearing off the roof of the car and killing the human occupant. A supercomputer running automated trading systems raises a value of a portfolio to US\$2.5 billion, then racks up losses resulting in a net loss of US\$22 million dollars. These stories are not the stuff of fearmongering technophobes, but actual incidents with very real legal and financial implications, quite aside from the cost of human life.

In the first case, an operator of a Tesla vehicle had activated the car’s Autopilot system and was travelling at over 100km/h. The system failed to detect a white semi-truck trailer and drove into it. This was the second incident of its kind.

In the second case, an investment fund MMWVWVW Limited (“VWM”) entered into a contract with a Monaco based investment firm, Tyndaris SAM (“Tyndaris”). Tyndaris promised an Artificial Intelligence (“AI”) managed investment account, running on a supercomputer capable of applying “machine learning” to market and social media data, to make emotionless, bias-free decisions. Tyndaris represented that the system had been extensively tested. While initially promising, it quickly lost staggering sums of money. VWM wanted out and claimed for the losses, while Tyndaris claimed for unpaid fees for the use of the system.

In dealing with any AI system, regardless of whether as end-user, retailer, or developer, there are a myriad snares to catch the unwary. This article aims to discuss the vexed question of legal liability when an AI goes awry, and what can be

done to protect legal interests when dealing with AI.

### IT TAKES A VILLAGE TO RAISE AN...AI?

AI does not spring fully formed from the head of some unwitting person. In its current implementations, it is mostly an advanced piece of software with gargantuan data processing power, taking in information and refining its internal parameters in order to achieve some predefined outcome. During the training phase, programmers define a desired outcome, and provide the AI data that “is” or “is not” the desired outcome. The AI considers this data, and (hopefully) develops a system to effectively discriminate between the positive and negative outcome.

An AI system intended to drive cars is trained on vast amounts of traffic data. For example, the AI is programmed to recognize that a traffic light showing red means “stop”. All kinds of images are fed into the AI and some of these are traffic lights showing red. The AI then “learns” that images with certain characteristics are “red lights”, which mean “stop”. When it “sees” a new traffic light on the road, the AI compares it to its understanding of a “red light”, and determines if it should stop or go. The programmer does not set rules such as ‘red lights look like this but not that’, the AI develops the rules itself.

In the same manner, an AI intended to run a stock portfolio would be given data about the market, and it would identify characteristics that presage a rise or a fall in the stock market from infinitely more data points than a human fund manager could.

However, there are multiple actors in this transaction. There is the programmer, who wrote the program to begin with. There is the provider of the data, which often may be a different entity. There is the company who employed the programmer. Where the AI is sold as a product or service, there is the end-user of the AI. As with all transactions, everything is fine, until it is not. The question then arises as to who is at fault, and where liability should lie.

### RISK OF (BLAMING) THE MACHINES?

Before going any further, it should be said that the idea of affixing liability onto the AI itself is an exercise in futility. AI does not have legal personality. Liability can only fall on the individuals or legal entities such as companies that among others, create, develop, or interact with it.

## Financial systems

The use of computerized systems to conduct trades is not a new one. IBM showed as far back as 2001 that computerized systems could outperform human traders.<sup>1</sup> In the Singapore case of *B2C2 v Quinone*, the Singapore International Commercial Court had the opportunity to comment on the issue of automated trading systems, which would execute trades on cryptocurrency exchanges when certain predefined parameters were met.

The court held that for the operation of the contractual doctrine of mistake as a result of trades conducted by the automated systems, the relevant state of mind was that of the programmer of the system.

However, the system in *B2C2* was deterministic, i.e. the decision made by the system was fully predictable given knowledge of the initial inputs. The rules of the system were all coded in by the programmer, and it would not be inaccurate to impute her knowledge to that of the system. This is in contrast to AI. Once the system is set in motion and trained, the AI can make decisions based on opaque and esoteric logic known only to it. Affixing liability on the programmer, quite aside from questions of privity of contract, for the workings of a system she may not be able to explain, trained on data which she may not have provided, is a vastly unsatisfactory solution.

## Physical systems

The field of self-driving vehicles has some guidance in the form of the Road Traffic (Autonomous Motor Vehicles) Rules 2017, which provides that a person who wishes to conduct a trial of or use an autonomous motor vehicle<sup>2</sup> must make an application to the Land Transport Authority<sup>3</sup> and take out liability insurance.<sup>4</sup>

This makes it clearer, in the event of an accident, how an injured party can seek redress – they would go after the person who made the application to use the vehicle. Better yet, they could claim against insurance.

<sup>1</sup> Agent-Human Interactions in the Continuous Double Auction <  
<http://spider.sci.brooklyn.cuny.edu/~parsons/courses/840-spring-2005/notes/das.pdf>>.

<sup>2</sup> Defined in s 2 of the Road Traffic Act as a vehicle with a system that enables the operation of the motor vehicle without the active physical control of, or monitoring by, a human operator

<sup>3</sup> Rules 5 and 6 of the Road Traffic (Autonomous Motor Vehicle) Rules 2017.

<sup>4</sup> Rule 14 of the Road Traffic (Autonomous Motor Vehicle) Rules 2017.

However, this is not the end of the question; there are use cases for AI in the physical world outside of autonomous vehicles. Companies may wish to use robots to do warehousing, a la Amazon and Alibaba. After all, robots, unlike humans, do not go on strike or demand dental benefits. There is therefore a question of affixing liability when things go wrong. In an example, an automated Amazon robot accidentally punctured a can of bear repellent at a warehouse, requiring 24 human employees to be sent to the hospital. Should these employees seek redress against Amazon for failing to provide a safe working environment, or should they instead sue the developers of the AI for having provided defective software? Is the developer of the AI the employee or contractor of Amazon such that rules excluding liability for injuries are caused by co-workers apply?

## PRACTICAL STEPS

While it may be tempting to wait for a definitive pronouncement on the liability of parties, any dispute litigated would likely be expensive and protracted. The Personal Data Protection Commission has suggested in its Proposed Model AI Governance Framework that AI decisions should be explainable, transparent and fair. The question is how this should be implemented in practice. In any event, the very nature of AI is that the decisions are not often explicable or transparent.

Until you can sue Skynet (or its developer, user, or employee), what do we do? The short answer is:- agree on risk allocation. The contractual frameworks for allocating risk in uncertain situations are well established, and provide more certainty than leaving liability as a question to be determined when the relationship sours. By clearly allocating responsibility and risk ahead of time, parties are also better advised as to how they can manage the use of the AI, and who must take efforts to exercise oversight over its operation. It may also be desirable to seek indemnities or limitations of liability, depending on whether one is an end-user or developer of the AI.

Further, there is also a question of ownership of data. As mentioned earlier, an AI as currently utilized learns through processing data. This means that the AI could (and would likely) be learning from the actual outcomes of each and every trade it makes. The contract should thus state whether the AI is allowed to do so, and if so, whether the developer is allowed to use these outcomes to improve the AI trading software.

Users and developers alike should be alive to the dangers inherent in the development and use of AI, and take care to ascertain the scope of their exposure and liability before embarking on any

grand plans involving AI. Its unique nature demands unique sandboxing of legal liabilities, but risk management can be best achieved in the short term through contractual negotiations between contracting parties.

---

If you have any questions or comments on this article, please contact:



**Rakesh Kirpalani**

Director, Dispute Resolution & Information  
Technology  
Chief Technology Officer  
T: +65 6531 2521  
E: rakesh.kirpalani@drewnapier.com

**Timothy Oen**

Associate, Dispute Resolution & Information  
Technology  
E: timothy.oen@drewnapier.com

---

**Updates in DrewTech Series**

1. Chapter 1: The Importance of An Exit Strategy in Technology Contracts <6 March 2019>
2. Chapter 2: Employees, Technology and A Legal Hangover – Bring Your Own Problems? <4 June 2019>
3. Chapter 3: I host, you post, I get sued? <24 September 2019>
4. Chapter 4: Diabolus ex machina <18 February 2020>

*The content of this article does not constitute legal advice and should not be relied on as such. Specific advice should be sought about your specific circumstances. Copyright in this publication is owned by Drew & Napier LLC. This publication may not be reproduced or transmitted in any form or by any means, in whole or in part, without prior written approval.*

Drew & Napier LLC  
10 Collyer Quay  
#10-01 Ocean Financial Centre  
Singapore 049315

[www.drewnapier.com](http://www.drewnapier.com)

T : +65 6535 0733  
T : +65 9726 0573 (After Hours)  
F : +65 6535 4906