

AM I AN ALGORITHM OR A PRODUCT?
WHEN PRODUCTS LIABILITY SHOULD APPLY TO ALGORITHMIC
DECISION-MAKERS

Abstract

Over the years mankind has come to rely increasingly on machines. Technology is ever advancing, and in addition to relinquishing physical and mere computational tasks to machines, algorithms' self-learning abilities now enable us to entrust machines with professional decisions, for instance, in the fields of law, medicine and accounting.

A growing number of scholars and entities now acknowledge that whenever certain "sophisticated" or "autonomous" decision-making systems cause damage, they should no longer be subject to products liability but deserve different treatment from their "traditional" predecessors. What is it that separates "traditional" algorithms and machines that for decades have been subject to traditional product liability legal framework from what I would call "thinking algorithms," that seem to warrant their own custom-made treatment? Why have "auto-pilots," for example, been traditionally treated as "products," while autonomous vehicles are suddenly perceived as a more "human-like" system that requires different treatment? Where is the line between machines drawn?

Scholars who touch on this question, have generally referred to the system's level of autonomy as a classifier between traditional products and systems incompatible with products liability laws (whether autonomy was mentioned expressly, or reflected in the specific questions posed). This article, however, argues that a classifier based on autonomy level is not a good one, given its excessive complexity, the vague classification process it dictates, the inconsistent results it might lead to, and the fact said results mainly shed light on the system's level of autonomy, but not on its compatibility with products liability laws.

This article therefore proposes a new approach to distinguishing traditional products from "thinking algorithms" for the determining whether products liability should apply. Instead of examining the vague concept of "autonomy," the article analyzes the system's specific features and examines whether they promote or hinder the rationales behind the products liability legal framework. The article thus offers a novel, practical method for decision-makers wanting to decide when products liability should continue to apply to "sophisticated" systems and when it should not.

AM I AN ALGORITHM OR A PRODUCT? WHEN PRODUCTS LIABILITY SHOULD APPLY TO ALGORITHMIC DECISION-MAKERS

*Karni A. Chagal-Feferkorn**

INTRODUCTION.....	2
PART I: WHY NOT USE "AUTONOMY" TO DISTINGUISH PRODUCTS FROM THINKING ALGORITHMS	10
<i>A. On Automation and Autonomy</i>	10
<i>B. The Very Many Faces of Autonomy</i>	13
PART II: PRODUCTS LIABILITY RATIONALES.....	19
PART III: IS IT A PRODUCT? IS THERE A DEFECT?	26
<i>A. On "Products".....</i>	26
<i>B. On "Defects".....</i>	28
PART IV: DISTINGUISHING PRODUCTS FROM THINKING ALGORITHMS.....	31
<i>A. Examples of Products versus Thinking Algorithms.....</i>	31
<i>B. Which Products Liability Rationales Apply to Which Algorithms.....</i>	35
1. Promoting Safety.....	35
2. Avoiding a Chilling Effect.....	49
3. Victims' Compensation.....	51
<i>C. Keep it Simple, System.....</i>	54
CONCLUSIONS.....	56

** Fellow, Haifa Center for Law & Technology, University of Haifa Faculty of Law. I thank Niva Elkin-Koren for her endless support that made this research possible. I would also like to thank Yonatan Arbel, Michal Gal, Yaron David, Itiel Dror, Orr Dunkelman, Rita Osadchy, Amnon Reichman, Catherine Sharkey and Avinoam Zimmerman. This research was supported by the Center for Cyber Law & Policy (CCLP), established by the University of Haifa in collaboration with the Israeli National Cyber Bureau. Any mistakes or omissions are the author's.

INTRODUCTION

Culminating in the Industrial Revolution, machines and tools replaced or assisted humans in performing physical tasks.¹ Drills, engravers, weaving machines and the like employed machines' physical advantages to free human beings from Sisyphean physical labor. In general, damage caused by such tools has been governed, since the beginning of the 20th century, by the legal framework of Products Liability.² Under this system the seller, manufacturer, distributor, or any other party in the distribution chain of a defective product is liable for the physical harm caused to the user or her property.³ Thus, victims of tools that have caught fire or came apart,⁴ of motor vehicles that have crashed,⁵ of food containing external undesired objects,⁶ etc., could all have brought successful products liability claims against the products' manufacturers or sellers.

¹ Ronald M. Hartwell, *THE INDUSTRIAL REVOLUTION AND ECONOMIC GROWTH* (Methuen & Co. LTD, 1971), 295-297; Nick von Tunzelmann, *Time-Saving Technical Change: The Cotton Industry in the English Industrial Revolution*, 32 *Explorations in Economic History*, 1-27 (1995). See also Steven King and Geoffrey Timmins, *MAKING SENSE OF THE INDUSTRIAL REVOLUTION; ENGLISH ECONOMY AND SOCIETY 1700-1850* (Manchester University Press, 2001), 68-97 (arguing that adoption of technological measures versus "hand technologies" was associated with more difficulties and inefficiencies than commonly perceived).

² Donald G. Gifford, *Technological Triggers to Tort Revolutions: Steam Locomotives, Autonomous Vehicles, and Accident Compensation*, 11 *J. Tort Law* (forthcoming, 2018), 49-50. For an overview of products liability history see, for example, Jane Stapleton, *PRODUCT LIABILITY* (BUTTERWORTH & CO., 1994), 9- 29. Further discussion on the development of the products liability framework will be held in Part II.

³ RESTATEMENT (SECOND) OF TORTS § 402A (1965).

⁴ See, e.g. *White Consolidated Industry, Inc. v. Swiney*, 376 S.E.2d 283 (Va. 1989) (finding an electric stove that caught fire to be defected); *McKenzie v. SK Hand Tool Corp.*, 272 III App. 3d 451, 650 (N.E.2d 612 (1995) (Finding a ratchet wrench that came apart during usage to be defective).

⁵ See, e.g. *Gray v. Lockheed Aeronautical Systems*, 125 F.3d 1371 (11th Cir. 1997) (holding that a jet aircraft that crashed into the sea was defective); *Four Corners Helicopters v. Turbomeca. S.A.*, 979 F.2d 1434 (10th Cir. 1992) (applying strict liability on the engine manufacturer of a helicopter that crashed); *Collazo-Santiago v. Toyota Motor Corp.*, 149 F.3d 23 (1st Cir. 1998) (holding Toyota's design of its air bags was defective).

⁶ See, e.g. 7 Cf. *Coulter v. Am. Bakeries Co.*, 530 So. 2d 1009 (Fla. Dist. Ct. App. 1988) (affirming verdict for plaintiff who swallowed metal wire found in a doughnut); *Coca-Cola Bottling Co. v. Parker*, 451 So. 2d 786 (Ala. 1994) (affirming verdict for plaintiff who swallowed particles of glass found in a soft drink bottle).

Over time, and with the advance of technology, the "products" used by humans became more sophisticated. Machines were no longer used merely to replace humans in performing physical tasks. Instead their superior computational abilities were utilized to assist or replace humans in processing data. The electronic calculator, for example, allowed engineers, merchants, accountants and other professionals to provide better and quicker outputs;⁷ autopilots were installed in airplanes to improve flight safety through an automated system capable of processing huge amounts of information in split seconds,⁸ while cruise-control and auto-parking devices were similarly installed in cars.⁹ Despite the increasing level of "sophistication" of these machines and devices, manufacturers and sellers of these products were generally held to the traditional products liability legal framework,¹⁰ akin to the liability applied to the more "simple" or "less sophisticated" products.¹¹

⁷ *Electronic Calculators- Handheld*, The National Museum of American History, available at: <http://americanhistory.si.edu/collections/object-groups/handheld-electronic-calculators>; Nick Valentine, *The History of the Calculator*, The Calculator Site (Mar 24, 2014), available at: <https://www.thecalculatorsite.com/articles/units/history-of-the-calculator.php>

⁸ Tracy Irons-Georges, *ENCYCLOPEDIA OF FIGHT* (Salem Press, 2002), 120-122; Gangadhar B Kallur, *History of Autopilot*, available at: http://www.academia.edu/27170918/HISTORY_OF_AUTOPILOT; Kyle Colonna, *Autonomous Cars and Tort Liability*, 4 W. Res. J. L. Tech & Int. 81, 93-97 (2012).

⁹ *Ralph Teetor and the History of Cruise Control*, American Safety Council, available at: <http://blog.americansafetycouncil.com/history-of-cruise-control-2/>; *A Brief History of Cruise Control* (Jul 20, 2016), available at: <https://www.frontierchrysler.ca/history-cruise-control/>; Stephane Levis, *Look, Ma, No Hands: Self-Parking Cars Take the Wheel*, GEICO Insurance (May 5, 2016), available at: <https://www.geico.com/more/driving/auto/car-safety-insurance/look-ma-no-hands-parking-technology-takes-the-wheel/>.

¹⁰ See, e.g., *Richardson v. Bombardier Inc.* (M.D. Fla. 2005) where plaintiffs have (unsuccessfully) raised, among other things, manufacture defect claims pertaining to the autopilot system of an airplane that had crashed. Similar products liability claims were also raised in connection with a plane-crash in *Moe v. Avions Marcel Dassault-Breguet Avion*, 727 F.2d 917 (10th Cir. 1984). This time, defect claims against the autopilot system were upheld. For additional examples see Colonna, *Supra* Note 8.

¹¹ As will be elaborated on below, products liability may be imposed on the basis of "strict liability," requiring only the showing of a "defect" of the product, regardless of fault. Other cases require negligence by the defendant for liability to apply. The

Technology is ever advancing, and in addition to relinquishing physical and computational tasks to machines, algorithms' self-learning abilities now allow them to reach their own conclusions based on databases of previous cases.¹² This in turn enables us to entrust machines with making complex decisions that until lately required human discretion, and even replacing professional human judgment in matters of expertise where there is no clear "right or wrong" answer.¹³

In the field of Law, for example, "virtual attorneys" such as IBM's "Ross" are utilized by law firms in conducting legal research,¹⁴ algorithmic ODR mechanisms solve disputes online, often without any human facilitator,¹⁵ and bail algorithms determine whether defendants awaiting trial may post bail to be released.¹⁶ Physicians too rely more and more on algorithms in order to diagnose medical conditions and select optimal treatments,¹⁷ while tax advisors, company directors and even priests are being significantly assisted, or even replaced, by algorithms.¹⁸

type of products involved may affect the specific type of products liability applied. See Part II for a more detailed discussion.

¹² See further discussion on how "self-learning" works in Part IV.

¹³ As will be discussed in more detail, such decisions are based on probabilities, such that only in hindsight does it become clear whether the decision was the right one or not.

¹⁴ John Mannes, *ROSS Intelligence Lands \$8.7M Series A to Speed Up Legal Research with AI*, Techcrunch (Oct 11, 2017), available at: <https://techcrunch.com/2017/10/11/ross-intelligence-lands-8-7m-series-a-to-speed-up-legal-research-with-ai/>; Anthony Sills, *Ross and Watson tackle the law*, IBM Watson Blog (Jan 14, 2016), available at <https://www.ibm.com/blogs/watson/2016/01/ross-and-watson-tackle-the-law>.

¹⁵ Michael Legg, *The Future of Dispute Resolution: Online ADR and Online Courts*, (July 18, 2016), forthcoming – Australasian Dispute Resolution Journal, available at SSRN: <https://ssrn.com/abstract=2848097>.

¹⁶ A. J. Wang, *Procedural Justice and Risk-Assessment Algorithms* (Jun 21, 2018), available at SSRN: <https://ssrn.com/abstract=3170136>.

Tom Simonite, *How to Upgrade Judges with Machine Learning*, MIT Tech. Rev., (March 6, 2017).

¹⁷ See, e.g., Alina Shrourou, *Deep Learning in Healthcare: a Move Towards Algorithmic Doctors*, New-Medical (Mar 15, 2017), available at: <https://www.news-medical.net/news/20170315/Deep-learning-in-healthcare-a-move-towards-algorithmic-doctors.aspx>

Vinod Khosla, *Technology Will Replace 80% of What Doctors Do*, Fortune Magazine (Dec 4, 2012);

¹⁸ Richard Susskind & Daniel Susskind, *Technology Will Replace Many Doctors, Lawyers, and Other Professionals*, Harv. Bus. Rev. (2016). See also McKinsey's forecast on the percentage of actions currently performed by human professionals to be replaced by atomization. MCKINSEY GLOBAL INSTITUTE, AUTOMATION

Yet advanced as such machines and algorithms may be, occasionally they are still bound to cause damage. A futuristic transition to a world of robo-doctors, for example, could never achieve perfect rates of recovery by all patients, and some patients' condition would inevitably deteriorate as a result of decisions made by algorithms. Litigants relying on strategic advice provided by IBM's aforementioned "Ross" and the like are also bound to lose cases or negotiations from time to time, while some of the individuals filing tax reports prepared by a virtual tax advisor would certainly still be subject to tax investigations and sanctions due to the algorithm's choices. Indeed, where such decision-makers have already been put to test - in the field of driving - we have already witnessed that advanced algorithms are not immune to reaching erroneous decisions and causing damage.¹⁹

Should the traditional products liability framework continue to apply to the "new generation" of "decision-making" tools,²⁰ those that

POTENTIAL AND WAGES FOR US JOBS (2017) available at: <https://public.tableau.com/profile/mckinsey.analytics#!/vizhome/AutomationandUSjobs/Technicalpotentialforautomation>.

¹⁹ In March 2018, a self-driving Uber car in autonomous mode hit and killed a woman in Arizona. Sam Levin and Julia C. Eong, *Self-driving Uber Kills Arizona Woman in First Fatal Crash Involving Pedestrian*, The Guardian (Mar 19, 2018), available at: <https://www.theguardian.com/technology/2018/mar/19/uber-self-driving-car-kills-woman-arizona-tempe>. Less than a week later, a Tesla vehicle running on autopilot crashed and killed its driver. *Tesla Car that Crashed and Killed Driver was Running on Autopilot, Firm Says*, The Guardian (Mar 31, 2018), available at: <https://www.theguardian.com/technology/2018/mar/31/tesla-car-crash-autopilot-mountain-view>. In May 2016, a Tesla driver died as the vehicle's sensors failed to detect a trailer crossing the highway (Peter Valdes Dapena, *Tesla in Autopilot Mode Crashes into Fire Truck*, CCN-tech (Jan 24, 2018); *Tesla Driver Dies in First Fatal Crash while Using Autopilot Mode*, The Guardian (Jun 30, 2016)). While the accident rate of autonomous vehicles is expected to drop significantly, accidents and fatalities cannot be avoided completely. See, e.g., John Villasenor, *Products Liability and Driverless Cars: Issues and Guiding Principles for Legislation*, Brookings (Apr 24, 2014), available at <https://www.brookings.edu/research/products-liability-and-driverless-cars-issues-and-guiding-principles-for-legislation/>; Bryant W. Smith, *Automated Driving and Product Liability* 1 Mich. St. L. Rev. 7-20 (2017).

²⁰ This article would not distinguish between robotic decision makers or algorithmic decision makers that have no physical embodiment. Rather, it would use phrases

replace human discretion and enjoy rising levels of "autonomy" and self-learning abilities?²¹

Several scholars have argued that certain "sophisticated" or "autonomous" decision-makers require treatment different from their "traditional" predecessors. In other words, while technology is always advancing, many believe that there is something "unique" in contemporary self-learning algorithms that separates them from the general group of "machines" or "products," and warrants a tort legal framework different from products liability that would apply to damages caused by them.

The European Parliament, for example, has issued a draft report explaining that ordinary rules on liability are insufficient for

such as "sophisticated systems"; "self-learning algorithms" or "autonomous robots" interchangeably. See Jack M. Balkin, *The Three Laws of Robotics in the Age of Big Data*, 78 Ohio St. L.J. (2017) (suggesting that both algorithms and robots are similar members of the "algorithmic society"). For further discussion on the decision to treat both algorithms and robots alike for the purpose of discussing applicable liability framework see Karni Chagal-Feferkorn, *The Reasonable Algorithm*, 1 U. Ill. J. L. Tech. & Pol'y 116-117 (2018).

²¹ This question is of even greater importance due to the shift from "human afflicted damages" to "algorithmic afflicted damages" (caused precisely because machines now replace humans for increasingly more tasks) which enlarges the number of instances where said issue is expected to be invoked. In other words, technology now offers to shift decisions and actions from humans to machines, leading to a much larger share of cases that would potentially involve products liability claims. Smith, *Supra* Note 19. First, in certain fields technology dispenses with any human involvement. For example, while in the past damages associated with cleaning the house were the fault of a cleaning-person, the shift to cleaning robots will now invoke claims against a "product" or a "machine" once the cleaning has caused damage (such as the case where a cleaning robot in South Korea 'sucked' up its owner's hair while she was sleeping on the floor, mistaking it for dirt. Justin McCurry, *South Korean Woman's Hair 'Eaten' by Robot Vacuum Cleaner as She Slept*, The Guardian (Feb 9, 2015), available at: <https://www.theguardian.com/world/2015/feb/09/south-korean-womans-hair-eaten-by-robot-vacuum-cleaner-as-she-slept>). See Colonna, *Supra* Note 8, 102-104. Secondly, even when humans remain "in the loop" and will merely be assisted by sophisticated algorithms, a larger share of the damaging decisions will be reached by algorithms rather than the person; thus again, the focus is shifted from human fault to products liability. Smith gives an example from the world of driving, where in the past damaging decisions often stemmed from a combination of both human and machine failures. For instance, car accidents involving some sort of manufacture defect were frequently caused or made worse because in addition to the car defect the human driver was speeding or was acting in some other form of faulty driving. Driverless vehicles, on the other hand, will themselves account for most or all real-time driving decisions, and any defect in their manufacturing or design will therefore play a greater - or a sole - role in future car accidents. *Ibid.*, 29-30

autonomous robots, since they can no longer be considered tools in the hands of other actors. The report suggested granting autonomous robots an independent legal status of “electronic persons,” which would even allow these robots themselves to pay damages for the harm they cause.²² Of the committees providing recommendations based on said report, the committees of Transport and Tourism, of Employment and Social Affairs and of Environment, Public Health and Food Safety have all agreed that new liability rules ought to be developed to account for the new character of robotic decision makers.²³

In the field of medicine, several scholars have proposed treating medical applications in a way more compatible with how human professionals are treated after causing damage rather than as medical devices. Bambauer, for example, suggested that certain medical

²² “[w]hereas the more autonomous robots are, the less they can be considered simple tools in the hands of other actors (such as the manufacturer, the owner, the user, etc.); whereas this, in turn, makes the ordinary rules on liability insufficient and calls for new rules which focus on how a machine can be held – partly or entirely – responsible for its acts or omissions...” 2015/2103 (INL) European Parliament Draft Report on Civil Law Rules on Robotics (May, 2016), Section S.

²³ “[w]hereas, for the purpose of civil liability, a distinction should be drawn between automated vehicles (incorporating devices allowing the automatic execution of some driving operations) and autonomous vehicles (which perform all such operations)... whereas in the former case the civil liability regime remains unchanged compared to that with conventional vehicles, while it needs to be adjusted in the latter case;” Opinion of the Committee on Transport and Tourism (Nov 16, 2016) with recommendations to the Commission on Civil Law Rules on Robotic (2015/2103(INL)); “[c]onsidering the increasing level of autonomy of robots, this should be accompanied by amending the rules on liability concerning the consequences associated with the actions or inaction of robots; is concerned by the lack of general framework and legal provisions with regard to work automation in this new and ongoing industrial revolution and considers it to be essential for the Union to specify a legal framework that reflects the complexity of robotics and its numerous social implications”; Opinion of the Committee on Employment and Social Affairs (Nov 9, 2016) with recommendations to the Commission on Civil Law Rules on Robotic (2015/2103(INL)); “[c]alls on the Commission and on the Member States to promote the development of assistive technologies, also through liability schemes that are different from those currently applicable, in order to facilitate the development and adoption of these technologies...” Opinion of the Committee on The Environment, Public Health and Food Safety (Oct 14, 2016) with recommendations to the Commission on Civil Law Rules on Robotic (2015/2103(INL)).

applications should not be regulated as medical devices, but more similarly to how human doctors are regulated.²⁴ Allain too compared IBM's Watson units used in the medical field to a "consulting physician,"²⁵ while Chung and Zink likened it to a "medical student," explaining that Watson was not a typical medical device and that the products liability regime would not suit it.²⁶ In the field of transportation, Abbott distinguished a conventional automobile from a driverless car which warrants a separate treatment of scrutinizing its "actions" as compared with those of a reasonable human driver.²⁷ In a previous paper I too argued that in general, algorithms replacing a human's professional judgment should be subject to the "reasonableness" standard that currently applies to humans rather than being treated as a product.²⁸

Employment of sophisticated machines that outperform people, however, is not new. Algorithms themselves date back more than 2,000 years.²⁹ Automated machines of different kinds were described in detail as early as 800 years ago.³⁰ What is it that separates "traditional" algorithms and machines, which for decades have been subject to the traditional products liability legal framework, from what I will generally refer to as "*thinking algorithms*" that seem to warrant their own custom-made treatment? Why have "auto-pilots," for example, been traditionally treated as "products"³¹ while autonomous vehicles are suddenly seen as a more "human-like" system that

²⁴ Jane R. Bambauer, *Dr. Robot*, 51 UC Davis Law Review 383 (2017).

²⁵ Jessica S. Allain, *From Jeopardy! to Jaundice: The Medical Liability Implications of Dr. Watson and Other Artificial Intelligence Systems*, 73 La. L. Rev. (2013).

²⁶ Jason Chung and Amanda Zink, *Hey Watson, Can I Sue You for malpractice? Examining the Liability of Artificial Intelligence in Medicine* (November 23, 2017). Forthcoming, *Asia-Pacific Journal of Health Law, Policy and Ethics*. Available at SSRN: <https://ssrn.com/abstract=3076576>.

²⁷ Ryan Abbott, *The Reasonable Computer: Disrupting the Paradigm of Tort Liability*, 86 Geo. Wash. L. Rev. 1 (2017).

²⁸ Chagal-Feferkorn, *Supra* Note 20.

²⁹ The Euclidean algorithm, for example, was invented around 300 BC and is still in use in existing technologies. *The Euclidean Algorithm*, Khan Academy, available at <https://www.khanacademy.org/computing/computer-science/cryptography/modarithmetic/a/the-euclidean-algorithm>.

³⁰ Ibn al-Razzaz al Jazari, *THE BOOK OF KNOWLEDGE OF INGENIOUS MECHANICAL DEVICES* (1206).

³¹ See, for example, the case of the 2013 Asiana-Air crash in San-Francisco, where the underlying legal actions consisted of products liability claims raised against Boeing, the manufacturer of the autothrottle that allegedly failed. Matt Hamilton, *Asiana crash: 72 passengers settle lawsuits against airline*, L.A. Times (March 3, 2013), available at <http://www.latimes.com/local/lanow/la-me-ln-asiana-airlines-settle-lawsuits-20150303-story.html>.

requires different treatment? Where is the fine line drawn between "products" and "decision-makers"?

While several scholars have touched on distinguishing "traditional" from "sophisticated" technologies for the purpose of applying products liability, no in-depth discussion specifically on this question has yet been held. Moreover, all potential distinguishing parameters that have been mentioned relate to the system's '*level of autonomy*' (while some referred to 'autonomy' expressly,³² other scholars discussed whether the system is able to wholly replace humans³³ or whether it outperforms humans³⁴ all aspects are associated with the system's level of autonomy, as will be discussed below). This article, however, demonstrates that 'autonomy' is not a desirable classifier between products and thinking algorithms, given its excessive complexity, the likelihood said classifier will yield absurd or inconsistent results, and the non-practical nature of the results obtained using autonomy as a classifier. A Roomba vacuum robot, for example, does replace a person in its action of cleaning; it (arguably) does so in a manner that outperforms the human cleaner,³⁵ and it possesses sufficient levels of autonomy that enable it to "decide" on its own in which direction to turn in order to continue cleaning. Yet the Roomba is probably not a "thinking algorithm" in the sense of requiring a special tort legal framework that would apply to it rather than the traditional products liability framework.

This article therefore proposes a new approach for distinguishing traditional products from "thinking algorithms" for the purpose of determining whether products liability should apply. Instead of examining the system's characteristics in isolation, I propose a "purposive interpretation" approach: one that analyzes the system's characteristics vis-à-vis the rationales behind the products liability legal framework, and identifies those that promote said rationales versus ones adversely affecting them. The article would thus offer a

³² European Parliament Draft Report on Civil Law Rules on Robotics, *Supra* note 22.

³³ Abbott, *Supra* Note 27.

³⁴ *Ibid.* See also Jason Millar and Ian Kerr, *Delegation, Relinquishment and Responsibility: The Prospect of Expert Robots*, in *ROBOT LAW* 102 (Elgar, 2016).

³⁵ Though not of statistical significance, it certainly outperforms the author of this article.

novel, practical method for differentiating "traditional products" from "thinking algorithms," based on fulfilling the rationales behind products liability laws and hence provide decision-makers with tools to better decide when products liability should apply and when it should not.

Part I explores the concept of autonomy in algorithmic decision-makers and shows why the different aspects of autonomy are insufficient for determining which systems should and which should not be subject to products liability laws. Part II provides background to the legal framework of products liability and its rationales. Part III then addresses a more preliminary discussion of whether "thinking algorithms" even fall within the realm of "products liability," given that they do not necessarily reflect a "product" and that damaging decisions reached by them are not necessarily the result of a "defect." Part IV analyzes how different characteristics of sophisticated systems affect the rationales behind the products liability framework. It thereby points out the characteristics that do not promote said rationales and should therefore indicate that the underlying system warrants a different legal treatment. Part V mentions alternative suggestions to imposing liability on damages caused by advanced technology and concludes.

PART I: WHY NOT USE "AUTONOMY" TO DISTINGUISH PRODUCTS FROM THINKING ALGORITHMS?

A. *On Automation and Autonomy*

Automated machines, characterized by different levels of sophistication, have been utilized by humankind for centuries.³⁶ In addition to "open loop control" systems, which executed automated tasks on an injective predetermined and unchangeable trajectory,³⁷ "closed loop control" systems have for centuries been capable of automatically choosing among predetermined options, based on "real-time" feedback. "Open loop control" windmills, for example, employed the power of wind to grind grain automatically whenever

³⁶ Ibn al-Razzaz al Jazari, *Supra* Note 30. The dictionary definition of "automation" is "the technique of making an apparatus, a process, or a system operate automatically" or as "the state of being operated automatically". The phrase "automatic" is then defined as "having a self-acting or self-regulating mechanism". Merriam-Webster.com <https://www.merriam-webster.com> (retrieved Aug 5, 2018).

³⁷ Otto Mayr, *THE ORIGINS OF FEEDBACK CONTROL* (The M.I.T. Press., 1970).

the wind blew. So long as the system was intact it would continue automatically to grind the same predetermined amount of grain on windy days, over and over again.³⁸ "Closed loop control" windmills, however, were also able to increase the output of the windmill by automatically enlarging the amount of grain poured into the mill as the wind blew stronger (based on a hopper suspended on ropes that received more "knocks" the stronger the wind blew, and automatically caused more grain to be dispensed).³⁹ Modern examples of closed loop automation systems are ubiquitous. They include Roomba vacuum robots (altering their routes on the ground based on feedback such as striking obstacles);⁴⁰ automated parking systems (basing their motion on feedback received by their sensors and cameras)⁴¹ or even automated snore stopper pillows (a microphone embedded in the pillow identifies sounds of snoring, and based on that feedback automatically inflates the pillow to change the snorer's posture).⁴² Countless other old and new closed-loop control systems allow us to enjoy the benefits of automation daily.⁴³ The mere fact that the system adjusts its actions based on external feedback, however, has existed for generations and in itself certainly does not warrant legal treatment

³⁸ *Ibid.* A boiler set by timer to turn on at a certain timing is an example of a modern open loop control mechanism, as the automation of turning the boiler on is independent of any variable other than the timing which was originally set.

³⁹ A boiler operating with a thermostat is a close loop control mechanism, as its continuous operation depends on feedback pertaining to the current water temperature.

⁴⁰ Julia Layton, *How Robotic Vacuums Work*, How-Stuff-Work?, available at: <https://electronics.howstuffworks.com/gadgets/home/robotic-vacuum3.htm>.

⁴¹ Laure McQuarrie, *The Valet Park4u is a Smart Self-Parking System from Valeo*, Trend-hunter (Jun 3, 2014), available at <https://www.trendhunter.com/trends/automatic-parking>.

⁴² Michael Plishka, *Pillow Adjusts Your Sleep Position*, Trend-hunter (Feb 7, 2014), available at <https://www.trendhunter.com/trends/pillow-snoring-automatic>

⁴³ To mention just a few amusing examples, automated doors for pets open automatically when receiving feedback indicating pet movement; Special utensils reduce eating speed of eating are available. Based on motion feedback received by the fork's sensors, the utensil vibrates whenever the user is eating too fast. Rahul Kalvapalle, *The Petwalk Pet Door is Both Pet-Friendly and Eco-Friendly*, Trend-hunter (Mar 14, 2014), available at: <https://www.trendhunter.com/trends/pet-door>; HAPILABS 103 HAPfork Bluetooth-Enabled Smart Fork, Amazon.com, available at: <https://www.amazon.com/HAPILABS-103-HAPIfork-Bluetooth-Enabled-Smart/dp/B00FRPCQ9Q>; respectively.

different from that of traditional products. What then might assist in making such a distinction?

Legal scholars engaged in products liability and sophisticated systems have suggested various directions to answer said question. The above report by the European Parliament, for example, proposed defining a "smart robot" as one whose "*autonomy*" is established by its interconnectivity with the environment (potentially through the use of sensors), and its ability to adapt its action to changes in it.⁴⁴ Millar and Kerr refer to "expert" robots, which would be classified as such when on average they consistently performed a well-defined set of tasks traditionally associated with human expertise – and do so better than the average human expert.⁴⁵ Abbott focuses on the system's ability to replace humans, more particularly on its ability to determine for itself how to complete tasks as set by humans.⁴⁶ Bambauer separates applications based on mere "measurement" from applications whose function is "knowledge-based";⁴⁷ Chung and Zink focus on the "higher level" of the system's duties, which require abilities of interpretation and analysis, while also distinguishing different systems according to their ability to make decisions as well as implement them.⁴⁸

As the following review shows, these diverse and seemingly widely different tests all relate to different aspects of the system's "autonomy," whether or not the term autonomy is stated expressly. The problem in classifying a system as a product or a thinking algorithm based on its level of autonomy, however, is the complexity

⁴⁴ European Parliament Draft Report on Civil Law Rules on Robotics, *Supra* Note 22.

⁴⁵ Millar and Kerr, *Supra* Note 34. Though said definition may also suit a weaving machine replacing a human weaver in the 19th century, Millar and Kerr also explained that in general, experts are not measured by their ability to follow known instructions; instead they prove their expertise by "filling in the blanks" so that an identical set of instructions is more likely to succeed when performed by an expert, and more likely to fail when performed by someone else.

⁴⁶ "What distinguishes an ordinary product from a computer tortfeasor in this system are the concepts of independence and control. Autonomous computers, robots, or machines are given tasks to complete, but they determine for themselves the means of completing those tasks." Interestingly, Abbott rejects the test of whether or not the machine's actions were foreseeable as distinguishing traditional from autonomous products. Abbott, *Supra* Note 27.

⁴⁷ "[M]easurement apps will have traditional instruments as their nearest conceptual neighbor, while knowledge apps emulate doctors or, perhaps, the patients' informal networks of health advisers." Bambauer, *Supra* Note 24.

⁴⁸ Chung and Zink, *Supra* Note 26.

of the term "autonomy," which is much greater than might be intuitively assumed.⁴⁹ Moreover, applying "autonomy level" or different aspects of it as a classifier might in many cases lead to absurd or inconsistent outcomes, and in any event it merely provides a "fuzzy" test whose results are not necessarily practical.

B. *The Very Many Faces of Autonomy*

Using "autonomy level" as means for distinguishing traditional products from thinking algorithms for the purpose of applying products liability laws is undesired. One of the reasons is that such a classifier would be difficult to implement, as its level of complexity is excessively high.

First, "autonomy" is a spectrum rather than a binary classification.⁵⁰

Secondly, various spectrums of autonomy exist, each focusing on completely different aspects. One common measurement of autonomy is the system's *freedom to act without human involvement* (or the allocation of decision-making power to humans or machine). Sheridan's spectrum, for example, offers ten levels of autonomy, the lowest being a situation where all processes are accomplished by a human being, without any machine assistance; the highest level is where the machine selects the desired courses of action but also executes them, not even informing the human of its choice and ignoring the human altogether.⁵¹ Adding even more to the complexity

⁴⁹ In fact, it was argued that any attempt to define "autonomy" would inevitably be based on controversial assumptions. Noel Sharkey, *Staying in the Loop: Human Supervisory Control of Weapons*, in *Autonomous Weapons Systems*, in LAW, ETHICS, POLICY (Nehal Bhuta et al. eds., Cambridge University Press, 2016). See also William C. Marra and Sonia K. McNeil, *Understanding "The Loop": Regulating the Next Generation of War Machines*, 36 Harv. J.L. & Pub. Pol'y, 18 (2013): "[A]utonomy is a complex concept with many components that cannot be captured simply through a distinction based on decision-making independence."

⁵⁰ *Ibid.*, 22-26.

⁵¹ In between are levels where the machine does not execute actions without human approval; executes only after the human has had time to veto the action; executes and then informs the human, based on the human's request or executes and then informs the human on the machine's own discretion of whether to inform or not. Raja Parasuraman et al., *A Model for Types and Levels of Human Interaction with Automation*, 30 IEEE Transactions on Systems, Man and Cybernetics- Part A: Systems and Humans (2000); Marra and McKneil *Supra* Note 49, 24. The full list

of said spectrum, is that even with identical systems employing identical technology, the division of tasks to those performed by humans and those "offloaded" to the system varies between particular scenarios.⁵² A GPS system, for example, merely assists the driver in navigation when driving in a familiar neighborhood. But once the driver finds herself in unfamiliar surroundings, the same GPS suddenly does much more than occasional assistance. Rather, the allocation of 'decision-making power' then changes such that the system's directions fully dictate the driver's actions.⁵³

A second method of assessing autonomy is based on the system's *ability to replace humans*.⁵⁴ This measurement is in itself

of Sheridan's spectrum levels is as follows: **Level 1:** The computer offers no assistance, human must do it all. **Level 2:** The computer offers a complete set of action alternatives, and **Level 3:** Narrows the selection down to a few, or **Level 4:** Suggests one, and **Level 5:** Executes that suggestion if the human approves, or **Level 6:** Allows the human a restricted time to veto before automatic execution, or **Level 7:** Executes automatically, then necessarily informs the human. **Level 8:** Informs the human after execution only if the human asks, or **Level 9:** Informs the human after execution if it, the computer, decides to do so. **Level 10:** The computer decides everything and acts autonomously, ignoring the human completely. *Ibid.*, 21.

⁵² Itiel E. Dror and Jennifer L. Mnookin, *The Use of Technology in Human Expert Domains: Challenges and Risks Arising from the Use of Automated Fingerprint Identification Systems in Forensic Science*, 9 *Law, Probability and Risk* (2010). "[T]hese technologies may distribute cognition in a number of ways: They can work to help and support the human expert; they can work alongside the human expert in a collaborative partnership; or the technology can play a more critical and significant role than the human."

⁵³ *Ibid.*

⁵⁴ See, for example, the Society of Automotive Engineers' (SAE) spectrum for autonomous vehicles. Level 0 of the SAE spectrum is no automation, where human drivers perform all tasks required for driving; Level 5 is full automation where the vehicle itself is capable of performing all driving functions under all circumstances. In between, automation (not autonomy) is used as a tool that increasingly assists the human driver. *Automated Vehicles for Safety*, NHTSA, available at: <https://www.nhtsa.gov/technology-innovation/automated-vehicles-safety>. The full list of the SAE's spectrum levels is as follows: **Level 0: No Automation.** Zero autonomy; the driver performs all driving tasks. **Level 1: Driver Assistance.** Vehicle is controlled by the driver, but some driving-assist features may be included in the vehicle's design. **Level 2: Partial Automation.** Vehicle has combined automated functions, like acceleration and steering, but driver must remain engaged with the driving task and monitor the environment at all times. **Level 3: Conditional Automation.** Driver is a necessity, but it is not required to monitor the environment. Driver must be ready to take control of the vehicle at all times with notice. **Level 4: High Automation.** Vehicle is capable of performing all driving functions under certain conditions. Driver may have the option to control the vehicle. **Level 5: Full Automation.** Vehicle is capable of performing all driving functions under all conditions. Driver may have the option to control the vehicle. SAE's highest level of

branched and complex, as several sub-analyses have been suggested in that context. Most sub-analyses focus on the system's ability to adapt to changing conditions, although the specific tests vary significantly. They include the relatively simple test of whether the machine is limited to choosing among pre-programmed options or is capable of choosing options that are not fully pre-programmed.⁵⁵ They also include more complex tests, namely if the following three "attributes of autonomy" are met: *frequency of human operator interaction; machine's tolerance for environmental uncertainty; level of assertiveness of the machine.*⁵⁶ Another approach focuses on whether the system is able to perform all the following types of activity: "*skill-driven*," "*rules-driven*" and "*knowledge driven*." (These concern which of the system's abilities are required in order to perform. The last addresses its ability to make decisions when the anticipated "rules" or circumstances change and the system has to come up with the optimal reaction to a scenario it was not "prepared" for).⁵⁷

automation therefore still grants the human driver discretion to maintain or take control over the vehicle.

⁵⁵ Kenneth Anderson and Matthew Waxman, *Law and Ethics for Autonomous Weapon Systems- Why a Ban Won't Work and How the Laws of War Can*, Jean Perkins Task Force on National Security and Law, available at: http://media.hoover.org/sites/default/files/documents/Anderson-Waxman_LawAndEthics_r2_FINAL.pdf

⁵⁶ According to Marra and McKneil, a machine is the more autonomous the less frequently a human operator must intervene and give her instructions; the more adaptability it shows in the face of scenarios it is not fully programmed to encounter; and the more it is able to change its operating plan in order to achieve its pre-programmed task, for instance, when the machine is "stuck." Marra and McKneil *Supra* Note 49, 18-22.

⁵⁷ In more detail, human activity may roughly be divided into three groups: "*skill-driven*" activities, namely the ability to accomplish physical tasks; "*rules-driven*" activities, namely the ability to comply with pre-determined rules; and "*knowledge driven*," namely the ability to make decisions when the rules previously mentioned are inadequate. Antonio Chialastri, *Automation in aviation*, in *AUTOMATION 79*, Florian Kongoli ed., 2012). An autopilot, for instance, is used for "*skill-driven*" activities (such as maintaining very precise altitude); for "*rules-driven*" activities (for instance, landing a plane at the correct angle in certain wind and whether conditions - based on instructions or rules predetermined for such conditions); and theoretically for "*Knowledge driven*" abilities (for instance, deciding where to land a plane when a malfunction occurs and existing rules will obviously lead to a crash. The famous "Miracle on the Hudson", where U.S. Airways pilot Sully Sullenberger

A third method for evaluating autonomy refers to a "stronger" or "substantial" measure of autonomy that will probably be developed more in the future, focusing on the system's own cognitive-awareness and real freedom of choice.⁵⁸ The spectrum proposed by the Air Force Research Lab (AFRL), for example, refers at the highest levels of autonomy to systems that are "cognizant."⁵⁹

Thirdly, to each of the different forked tests discussed above, another distinct dimension of analysis must be added: the specific stage of the machine's decision-making process.

According to Boyd, the human decision-making mechanism comprises four steps: "Observe"; "Orient"; "Decide"; "Act."⁶⁰ This cycle, known as the "OODA Loop," is applicable to machines as well. Sheridan's model for machine's information-processing, for example, includes the very similar stages of "Information Acquisition"; "Information Analysis"; "Decision Selection"; "Action Implementation."⁶¹ Naturally, as a machine becomes capable of performing more of the four steps on its own, the better the odds that it will achieve a higher "autonomy score" on the spectrum of whether it can replace humans or not.⁶² But in addition to this partial overlap

safely force-landed an airplane after two of its engines suddenly failed, is an example of human's superiority in "knowledge driven" decisions, given that pilot Sullenberger made the "right decision" as opposed to flight algorithms which would have seemingly led to a catastrophic crash. Clint Eastwood's movie "Sully" focuses on that exact point: Sully, 2016. See also Adam Smith, *The Miracle on The Hudson: How it Happened*, The Telegraph (Nov 22, 2016)).

⁵⁸ See, e.g., Eliav Lieblich and Eyal Benvenisti, *The Obligation to Exercise Discretion in Warfare: Why Autonomous Weapon Systems Are Unlawful* (August 13, 2014). *Autonomous Weapons Systems: Law, Ethics, Policy* 244 (Nehal Bhuta et al., eds., Cambridge University Press, 2016, forthcoming), available at SSRN: <https://ssrn.com/abstract=2479808>.

⁵⁹ At its lowest level, The AFRL's spectrum refers to a system that is remotely controlled by a human, or executes missions entirely pre-planned by humans. At its mid-level stages of autonomy, the system itself may respond to real-time events. Marra and McKneil *Supra* Note 49, 25.

The full list of the AFRL's spectrum levels is as follows: **Level 0:** Remotely piloted vehicle. **Level 1:** Execute pre-planned mission remotely. **Level 2:** Changeable mission. **Level 3:** Robust response to real time faults/events. **Level 4:** Fault/event adaptive vehicle. **Level 5:** Real-time multi-vehicle coordination. **Level 6:** Real-time multi-vehicle cooperation. **Level 7:** Battlespace knowledge. **Level 8:** Battlespace single cognizance. **Level 9:** Battlespace swarm cognizance. **Level 10:** Fully autonomous.

⁶⁰ Frans P.B. Osinga, *SCIENCE, STRATEGY AND WAR; THE STRATEGIC THEORY OF JOHN BOYD* (Routledge, 2007), 1-3.

⁶¹ Parasuraman et al., *Supra* Note 51; Marra and McKneil *Supra* Note 49, 11

⁶² (Albeit not necessarily in the context of whether or not it is good at adapting to new conditions, which is not necessarily related to being able to execute all four stages of the OODA Loop.)

with the measures of autonomy discussed above, the OODA Loop may be independently used as a separate dimension on each of the autonomy measures. For instance, the "information acquisition" stage of a certain machine or algorithm may be characterized by high levels of autonomy based on the above spectrums, while the same system's "decision selection" stage might involve very little autonomy. Also, the information acquisition stage performed by a military drone might include gathering data on potential targets without the need for any human involvement – thus the drone receives a "high autonomy score" for the first test of autonomy. It might do so even in the face of changing weather conditions or new disguise methods used by the potential targets – thus receiving a high score for the second spectrum of autonomy as well. At the same time, the decision-making stage for exactly the same drone may involve very little autonomy - at least with respect to the first spectrum - as a decision to hit a target will very likely require human authorization and not be executed at the drone's own discretion.⁶³

As demonstrated above, the different measurements of "autonomy" are numerous and complex, and at times overlap. Should decision-makers decide to rely on all these measurements when determining whether products liability laws ought to apply, they would have to develop a complex matrix, accounting for all the different aspects of autonomy discussed. Moreover, since the matrix will likely not be of a "one size fits all" type, decision-makers would have to decide in advance, for each specific case of for each product type, how much weight to give each measurement (certain measurements will likely be relaxed and others more scrutinized, based on the specific type of system, its capabilities and its potential usages). But the complexity of developing the matrix and later of applying it in each case is not the only disadvantage stemming from using a system's level of autonomy as a classifier distinguishing

⁶³ For discussions on the importance of "keeping a human in the firing loop" of autonomous weapons see, for example, Michael Schmitt and Jeffrey S. Thurnher, *'Out of the Loop': Autonomous Weapons Systems and the Law of Armed Conflict*, 4 Harvard National Security Journal (2013); Kenneth Anderson and Matthew Waxman: *Law and Ethics for Robot Soldiers*, Columbia Public Law Research Paper No. 12-313 (2012), Available at SSRN: <http://ssrn.com/abstract=2046375>; Markus Wagner, *Taking Humans Out of the Loop: Implications for International Humanitarian Law*, 21 J.L. Info. & Sci. (2011).

traditional products from thinking algorithms. Additional disadvantages include the fact that the measurements themselves are "fuzzy". For example, when determining the "*system's tolerance to environmental changes*", the decision maker would not be facing a "yes or no" question, but would have to come up with an out-of-context numerical or qualitative estimation of the system's tolerance of such conditions. Moreover, factoring in all these "fuzzy" estimations into a combined outcome that classifies the system's level of autonomy in turn too would lead to a mere general sensation of whether the system is autonomous or not. Should the outcome be situated "somewhere in the middle" of the scale rather near any of its ends, it will not be a useful tool for determining if we are in the traditional product kingdom or have found ourselves in the realm of thinking algorithms. In fact, a classifier based on the level of autonomy might lead to absurd results or ones incompatible with current classifications (such as classifying a Roomba vacuum cleaner as a thinking algorithm, as demonstrated in the introduction, using a few of the autonomy measurements later discussed in depth). Furthermore, the "score" or outcome received, even if placed on one of the ends of the autonomy scale, would merely give us an indication of the system's autonomy. It would not, however, give us any indication on the desirability of applying products liability laws to the system.

Choosing but a few autonomy measurements to rely on rather than the whole matrix might render the process less complicated, but would still maintain all the very many other disadvantages. Such an approach would also increase the risk of reaching absurd or inconsistent results.⁶⁴

This article therefore introduces a different approach to distinguishing products from thinking algorithms. Rather than focusing on a general out-of-context analysis of autonomy, the article analyzes how different "yes or no" features of traditional products and their self-learning counterparts reconcile with the rationales behind products liability laws- rationales which are the center of the discussion below.

⁶⁴ As demonstrated above, analyzing a drone's level of autonomy based on its ability to act without humans' involvement while leaving aside the underlying OODA Loop stage would be meaningless. Similarly, focusing on a system's ability to replace humans or outperform them might render a coffee machine or even a 19th century engraver "autonomous".

PART II: PRODUCTS LIABILITY RATIONALES

The birth of products liability, the most “popular” of all case types in the US,⁶⁵ is attributed to technological advances: the shift from local craftsmen to mass production factories caused a “lack of privity” problem that eliminated victims’ means of redress.⁶⁶ The introduction of products liability resolved that discrepancy, by eliminating the requirement of privity of contract between victim and defendant.⁶⁷ Under products liability laws the seller or manufacturer of a defective product in a condition that is unreasonably dangerous is liable for the physical harms caused to the user or her property, even when there is no contractual relationships between them.⁶⁸ One of the main rationales behind products liability therefore is *compensating the victim*, which stems from corrective justice principles, where the tortfeasor is required to correct the wrong she has committed, based on justice and fairness considerations.⁶⁹ Comprising the general

⁶⁵ Tens of thousands of products liability cases are filed annually, more than any other case type. Ronald C. Porter, Lex Machina 2018 Product Liability Litigation Report, March 2018. In 2017, for example, 42,789 products liability cases were filed in the US. The total number of Commercial, IP, employment, antitrust, securities and bankruptcy cases filed that year combined was lower than said figure.

⁶⁶ See, e.g., *Winterbottom v. Wright* (1842), 152 Eng. Rep. 402, where the English Court of Exchequer decided a case where a mail coach collapsed and injured the plaintiff. It held: “*There is no privity of contract between the parties; and if the plaintiff can sue, every passenger, or even any person passing along the road, who was injured by the upsetting of the coach, might bring a similar action. Unless we confine the operation of such contracts as this to the parties who entered into them, the most absurd and outrageous consequences, to which I can see no limit, will ensue.*”

⁶⁷ Gifford, *Supra* Note 2. In 1916, The New York Court of Appeals in *MacPherson v. Buick Motor Co.*, 111 N.E. 1050, held that Buick owed a duty of care to its users despite not having privity of contract with them: “*If to the element of danger there is added knowledge that the thing will be used by persons other than the purchaser, and used without new tests, then, irrespective of contract, the manufacturer of this thing of danger is under a duty to make it carefully.*”

⁶⁸ RESTATEMENT (SECOND) OF TORTS, § 402A (1965).

⁶⁹ In *Escola*, the Supreme Court of California held that: “*the cost of an injury and the loss of time or health may be an overwhelming misfortune to the person injured, and a needless one, for the risk of injury can be insured by the manufacturer.*” *Escola v. Coca Cola Bottling Co.*, 150 P.2d 436 (Cal. 1944). See also John C.P. Goldberg and Benjamin C. Zipursky, *The Easy Case for Products Liability Law: A Response to Professors Polinsky and Shavell*, 123 Harv. L.Rev. 1919, 1944 (2010): “*In fact, we think that the case for allowing persons injured by defective products to obtain*

rationale of compensating the victims – which is one of the rationales behind the general legal framework of tort law⁷⁰ – are sub-rationales that are more specific to product liability in particular. First, a manufacturer or seller marketing their product for public consumption must assume a "special responsibility" toward consumers who might be injured by the product.⁷¹ A related aspect is that consumers who bought products that caused injuries did so relying on manufacturers' advertisements regarding their products, including their safety.⁷² So it is only "fair" that when the sense of safety and security customers relied on turns out to be unjustified, they will be compensated.⁷³ Another sub-rationale relevant in the context of the fairness of compensating the victim is that manufacturers or sellers are generally in a better position to absorb or spread the costs of damages caused by their products,⁷⁴ or to insure against them.⁷⁵ Thus, rather than having the unfortunate victim shoulder the overwhelming weight of the entire damages she has suffered, the costs may be insured, allocated among the entire group of consumers (in the form of a price increase) or marked as "the costs of doing business" in manufacturers' budgets.⁷⁶

redress is very easy. It rests on the idea that a manufacturer bears a responsibility to avoid causing injury by sending a dangerously defective product into the stream of commerce and is supported by principles grounded in negligence and warranty."

⁷⁰ See, e.g., Ernest Weinrib, *LAW FOR LAW'S SAKE THE IDEA OF PRIVATE LAW* 1, 132 (Cambridge: Harvard University Press, 1995).

⁷¹ RESTATEMENT (SECOND) OF TORTS, § 402A (1965) cmt. c; Marshall S. Shapo, *THE LAW OF PRODUCTS LIABILITY* (Butterworth Legal Publisher, 1994).

⁷² John E. Montgomery and David Owen, *Reflections on the Theory and Administration of Strict Tort Liability for Defective Products*, 27 S.C.L.Rev. (1976).

⁷³ This is regardless of victims' means of redress in the form of breach of warranty, a contractual cause of action that in general does require privity between plaintiff and defendant. U.C.C. § 2-312 (2004); See e.g., *All West Electronics, Inc. v. M-B-W, Inc.*, 75 Cal.Rptr.2d 509 (Cal. Ct. App. 5th Cir. 6/4/98); *T.W.M. v. American Medical Systems, Inc.*, 886 F. Supp. 842 (N.D. Fla. 1995). The privity condition, however, is not required in all states. see, e.g., *Renaissance Leasing, LLC v. Vermeer Mfg. Co.*, 322 S.W.3d 112, 129 (Mo. banc 2010).

⁷⁴ Montgomery and Owen, *Supra* Note 72, 809-810. Escola, *Supra* Note 69.

⁷⁵ George L. Priest, *The Current Insurance Crisis and Modern Tort Law*, 96 Ya L.J., 1559 (1987); "[C]ourts justified third-party insurance coverage based on how easy it seemed to be for manufacturers or service providers to aggregate risks by adding an insurance premium to the price of the product or service".

⁷⁶ Montgomery and Owen, *Supra* Note 72.

A second main rationale behind tort law in general is "Deterrence."⁷⁷ In the context of product liability deterrence is translated into deterring manufacturers from creating dangerous products, that is, *promoting safety*. Naturally, the threat of liability encourages manufacturers to improve the safety of their products: the safer the product, the less likely it is to cause damage and the less likely are manufacturers to be sued and to pay damages.⁷⁸ Said rationale is indeed sensible in the context of product liability, given that manufacturers are best positioned to eliminate or reduce the risks associated with their products, considering that unlike consumers they possess information regarding the product and can ensure inspections and quality control measures.⁷⁹

As will be elaborated below, the safety promotion rationale could be absolute, in the sense that whenever a defect exists the manufacturer will be liable regardless of the precautions she has taken to minimize the damage threat. Such liability is classified as "strict liability," which requires no fault (and is mostly applied under "manufacture defect" scenarios, as will be discussed below).⁸⁰ Alternatively, the safety promotion rationale may be applied only to a limited extent, namely to the extent that it is feasible and economical

⁷⁷ Richard A. Posner, *The Value of Wealth: A Comment on Dworkin and Kronman*, 9 *LEG. STUD.* 243, 244 (1980); John C. P. Goldberg, *Twentieth Century Tort Theory*, 90 *Geo. L. J.* (2002).

⁷⁸ RESTATEMENT (THIRD) OF TORTS § 2 cmt. a (1998). "*On the premise that tort law serves the instrumental function of creating safety incentives, imposing strict liability on manufacturers for harm caused by manufacturing defects encourages greater investment in product safety...*". Strict liability for harm caused by manufacturing defects has been supported on the ground that it promotes investment in product safety. See, e.g., *Hoven v. Kelble*, 256 N.W.2d 379, 391 (Wis. 1977) ("*Strict liability is an effective deterrent; it deters the creation of unnecessary risks, or to put it positively, strict liability is an incentive to safety.*"); *U.S. Airways v. Elliott Equip. Co.*, No. 06-1481, 2008 WL 4425238, at 5 (E.D. Pa. Sept. 29, 2008). ("*[i]mposing strict liability here would serve as an incentive to safety because Fluidics appears to be involved in these types of contracts on a regular basis and is in a better position than a consumer to prevent circulation of defective products.*")

⁷⁹ Manufacturers possess information on the design and production process of the product, as well as on potential alternatives. The exception is the potential usages consumers might make of the product, which are not always in the scope of the manufacturers' knowledge or expectations. RESTATEMENT (THIRD) OF TORTS § 2(c) cmt. m.

⁸⁰ See discussion below.

for the manufacturer to take the extra precautions such that their benefits outweigh their costs⁸¹ (as will be elaborated below, this type of liability is reflected mostly in "design defect" scenarios).

While scholars debate which of these two liability frameworks will better fulfil the following sub-rationale,⁸² another utilitarian purpose for imposing product liability is to help ensure that products' prices reflect their true social cost. This should discourage excessive consumption of dangerous products (given that products' prices will rise the more dangerous they are, as a result of product liability litigation).⁸³

The above rationales, as well as the effects of technological advances, are also reflected in the specific mechanisms that govern the different kinds of defect in products, namely *manufacture defect*, *design defect* and *failure to warn*. Manufacture defect occurs when the product was not properly manufactured (for example, due to departure from the product's assembly specifications or use of non-appropriate materials). Under the laws of most states, manufacture defects expose sellers to strict liability, which does not require proof of any fault by the seller to be found liable, as long as the existence of a defect is proved.⁸⁴

Among other things, the introduction of strict products liability is attributed to technological progress. Increasingly complicated production processes render it difficult for victims to investigate the damaging product and to prove negligence in the manufacturing process. In consequence, the courts adopted the strict liability no-fault standard that allows victims' redress even as technology advances.⁸⁵

⁸¹ Naturally, which benefits and which costs should be considered have been the subject of much debate and interpretation. In *Cipollone v. Liggett Group, Inc.* 644 F.Supp. 283 (1986), for example, the US District Court determined that the benefits to be weighed do not include the benefits to the industry where the defendant operates or to its employees.

⁸² See, e.g., Mitchell A. Polinsky and Steven Shavell, *The Uneasy Case for Product Liability*, 123 Harv. L.Rev. 1436 (2010).

⁸³ *Ibid.*

⁸⁴ RESTATEMENT (THIRD) OF TORTS § 2(a) (1998) "[A] product contains a manufacturing defect when the product departs from its intended design even though all possible care was exercised in the preparation and marketing of the product"; John C. P. Goldberg and Benjamin C. Zipursky, *The Oxford Introductions to U.S. Law: Torts*, in THE OXFORD INTRODUCTIONS TO U.S. LAW (2010).

⁸⁵ Gifford, *Supra* Note 2, 50-51, citing Escola, *Supra* Note 69, where the Supreme Court held: "As handicrafts have been replaced by mass production with its great markets and transportation facilities, the close relationship between the producer

In addition to ensuring compensation for victims, strict liability - as explained above - strives to encourage manufacturers to improve the safety of their products. The rationale of safety promotion is also reflected in the second scenario of product liability, design defect, albeit in a more complex manner. As suggested by its title, a design defect is a flaw in the design of the product itself, regardless of how it is manufactured. Importantly, damages caused by algorithmic decision makers, which will be the main focus of our discussion in the following parts, are most likely to fall under the category of design defect rather than manufacture defect.⁸⁶

In the past design defect was ruled according to the "consumer expectations" test, that is, whether the product was dangerous beyond the expectations of an ordinary consumer.⁸⁷

Since no proof of negligence by the manufacturer or seller was required, the consumer expectations test could have been classified as a form of strict liability, reflecting again the rationales of compensating the victim as well as encouraging safety.⁸⁸ Nowadays however - again among other things for technology-related reasons that will be discussed below - the majority of states have moved away from the consumer expectations test to the risk utility test.⁸⁹ By that approach, a design defect occurs when foreseeable risks associated

and consumer of a product has been altered. Manufacturing processes ... are ordinarily either inaccessible to or beyond the ken of the general public. The consumer no longer has means or skill enough to investigate for himself the soundness of a product...."

⁸⁶ This is because problems associated with coding do not stem from occasional defects in manufacture but instead are implemented all along the product line, which is compatible with design defects. Patrick F. Hubbard, *'Sophisticated Robots': Balancing Liability, Regulation and Innovation*, 66 Fla. L. Rev. 1803, 1854 (2014).

⁸⁷ RESTATEMENT (SECOND) OF TORTS § 402A (1965); David G. Owen, *PRODUCTS LIABILITY LAW* (3d ed. 2014) 292-299.

⁸⁸ Gifford, *Supra* Note 2.

⁸⁹ Owen, *Supra* Note 87, 292-299; Jeffrey K. Gurney, *Sue My Car Not Me: Products Liability and Accidents Involving Autonomous Vehicles*, U. Ill. J. L. Tech. & Pol'y (2013). Other reasons for not adhering to the consumer expectations test are that it connotes contract-based rather than tort-based liability; that its application to bystanders- whom presumably had no expectations of the product- is problematic; and that obvious or patent defects might block victims means of redress as they could have expected the danger. Mary J. Davis, *Design Defect Liability: In Search of a Standard of Liability*, 39 Wayne L.Rev. 1217 (1993).

with the product could have been minimized by using a feasible safer alternative.⁹⁰ To go back to our rationale of safety promotion, the risk utility analysis strives to improve safety, but does so only to the extent such improvements are feasible and economical. This is because the risk utility test does not require the use of the safest design possible, but of the safest design whose costs will not exceed the safety benefits it contributes as compared with the alternatives.⁹¹ The calculations required to determine whether a safer alternative is or is not mandated are complex;⁹² still, manufacturers are deemed best positioned to perform them.⁹³

Asymmetrical information and manufacturers' better knowledge of the product and its potential risks are also a key factor in the third scenario invoking products liability, failure to warn. Under this doctrine manufacturers must adequately warn consumers of the existence of hidden dangers, as well as instruct them on the safe usage of the product, if they wish to avoid liability.⁹⁴ Under the Restatement, the requirement is for "reasonable" instructions or warnings, which largely subjects this type of defect to the negligence standard as well.⁹⁵

The general trend in product liability law has indeed shifted back toward a more negligence-based approach.⁹⁶ Although strict liability fulfils the rationales of victim's compensation and promotion of safety

⁹⁰ RESTATEMENT (THIRD) OF TORTS (1998).

⁹¹ Gurney, *Supra* Note 89, 236; David G. Owen, *Toward a Proper Test for Design Defectiveness: "Micro-Balancing" Costs and Benefits*, 75 Tex. L.REV. 1661 (1997). See also *Turner v. General Motors Corporation* 514 SW 2d 497 (1974) where a Texas Court of Appeals stated that: "If a change in design would add little to safety, render the vehicle ugly or inappropriate for its particular purpose, and add a small fortune to the purchase price, then a court should rule as a matter of law that the manufacturer has not created an unreasonable risk of harm".

⁹² (Taking into consideration the likelihood of damage, severity of damage, the costs of adopting a safer measure, the probability that the safer measure would indeed minimize the risk for damage or the magnitude of said damage and, in such cases, the degree to which the damage was indeed lower, etc.).

⁹³ *Supra* Note 79.

⁹⁴ RESTATEMENT (THIRD) OF TORTS, § 17 (1998), OWEN, *Supra* Note 87.

⁹⁵ RESTATEMENT (THIRD) OF TORTS, § 17 (1998) cmt. m.: "an overwhelming majority of jurisdictions supports the proposition that a manufacturer has a duty to warn only of risks that were known or should have been known to a reasonable person."

⁹⁶ Given that the majority of products litigations involve design and warning defects which are subject to the negligence standard, and given that design defects are now subject to the risk-utility test rather than the consumer expectations test in most states, which, as explained above, reflects a shift from a strict liability approach to the negligence approach. Gifford, *Supra* Note 2, 52-55

to a fuller extent,⁹⁷ other competing rationales and interests have swung the pendulum back to the fault-based negligence approach. First, increased likelihood of liability is generally expected to impede development and innovation, hence to create a "chilling effect" on technological advancement.⁹⁸ Secondly, increased levels of safety (potentially stemming from likelihood of liability), are expected adversely to affect different features of the product, including its pricing, ease of operation, appearance and additional factors related to consumers' preferences other than safety.⁹⁹ The legal framework of product liability therefore purports to strike an optimal balance between contradictory rationales and interests. Naturally, such a balance also depends on the specific type of product and of industry.¹⁰⁰ "Thinking" algorithms' unique features may well affect the optimal balance between rationales and interests. But before turning to said analysis, let us check whether such algorithms are at all relevant in the "product liability" context, or whether their features render the product liability legal framework completely inapplicable from the outset.

⁹⁷ Under strict liability, manufacturers pay for damages caused by defected products regardless of the level of care they have demonstrated. Victims are thus compensated a larger portion of the cases. At the same time, the cost of damage is internalized by manufacturers, such that the products' prices reflect their true risk. Polinsky and Shavell, *Supra Note 82*.

⁹⁸ See, e.g., Gifford, *Supra Note 2*, 52-55; Colonna, *Supra Note 8*, 109-111. See also a report on "roboLaw" by the European Union, concluding that products liability law might indeed create a chilling effect on the introduction of automated cars to the market. Final Report Summary - ROBOLAW (Regulating Emerging Robotic Technologies in Europe: Robotics facing Law and Ethics), European Commission Cordis, available at: http://cordis.europa.eu/result/rcn/161246_en.html.

⁹⁹ David G. Owen et al., *PRODUCTS LIABILITY AND SAFETY, CASES AND MATERIALS* (The Foundation Press Inc. 1996).

¹⁰⁰ See, e.g., Goldberg and Zipursky, *Supra Note 69*, rebutting Polinsky and Shavell's argument against the efficiency and desirability of products liability laws (Polinsky and Shavell, *Supra Note 82*) emphasizing how different industries are differently affected by such laws. See also Polinsky and Shavell's reply, stressing that their contention on "whether product liability is undesirable depends on the particular product". Mitchell A. Polinsky and Steven Shavell, *A Skeptical Attitude About Product Liability is Justified: A Reply to Professors Goldberg and Zipursky* Stanford Law and Economics Olin Working Paper No. 399 (2010), Available at SSRN: <https://ssrn.com/abstract=1689272>.

PART III: IS IT A PRODUCT? IS THERE A DEFECT?

One could argue that regardless of product liability rationales, algorithms that replace human discretion simply cannot be classified as "products." Alternatively, it could be argued that damages caused by such algorithms may not be attributed to "defects." If so, the argument goes, applying product liability rationales to distinguish which algorithms should be governed by product liability would simply be irrelevant, just like discussing free speech rationales in order to determine whether the howling of a wolf should be protected under the First Amendment would be irrelevant- as the underlying legal framework simply does not apply to begin with. The next section addresses these preliminary queries and explains why product liability rationales should nevertheless be the basis for analysis.

A. *On "Products"*

While several states have clearly defined the term "product" for the purpose of applying products liability,¹⁰¹ in general it is up to the courts to determine in each given case whether an underlying damaging object is indeed a "product."¹⁰² For many types of systems or machines that have caused damage the classification as "product" does not raise any question marks; but systems based on information - which naturally are the subject matter of our article - do fall within the "grey zone" of the "products kingdom" and occasionally have been excluded from it.

First, several courts held that information in itself did not constitute a "product" for the purpose of applying products liability, because it lacked tangible form.¹⁰³ Secondly, and focusing on the

¹⁰¹ Joseph L. Reutiman, *Defective Information: Should Information Be a "Product" Subject to Products Liability Claims*, 22 Cornell J.L. & Pub. Pol'y 181 (2012), referring, for example, to IDAHO CODE ANN. § 6-1402(3) (2008) (defining a product as "any object possessing intrinsic value, capable of delivery either as an assembled whole or as a component part or parts, and produced for introduction into trade or commerce. Human tissue and organs, including human blood and its components, are excluded from this term.");

¹⁰² *Ibid.* The Restatement of Torts (Third) which refers to the term "product" is drafted in a broad manner which seems to leave much flexibility as to which products fall within the definition and which do not. "[a] product is tangible personal property distributed commercially for use or consumption..." Michael D. Scott, *Tort Liability for Vendors of Insecure Software*, 67 Md. L.Rev. 426 (2008).

¹⁰³ *America Online, Inc. v. St. Paul Mercury Insurance Co.*, 242 F. Supp. 2d 459 (E.D. Va. 2002) "*the plain and ordinary meaning of the word tangible is something*

analogy of such systems to "professional services"¹⁰⁴ (an analogy even better suited to thinking algorithms which, as discussed above, replace human professionals), courts were sometimes inclined to rule that such services were not to be viewed as "products."¹⁰⁵

Yet in many other instances courts did treat information as a "product" and applied products liability laws when errors in the information caused damage. Considerations for treating information as a "product" were its integration with a physical object;¹⁰⁶ the informational object was mass-produced;¹⁰⁷ to an extent, it had dangerous potential.¹⁰⁸ Naturally, these considerations tend to exist in the "thinking algorithms" this article focuses on (given that thinking algorithms are often embedded in physical objects such as cellular phones or computers; that they are often mass-marketed; and that errors in them might cause deadly results).¹⁰⁹

that is capable of being touched or perceptible to the senses. Computer data, software and systems do not have or possess physical form and are therefore not tangible property as understood by the Policy."; Winter v. G.P. Putnam's Sons, 938 F.2d 1033, 1036 (9th Cir. 1991). (While the court in Winter acknowledged that the information at hand was embedded in a book, which is a tangible object, it nevertheless ruled that an encyclopedia of mushrooms which misled a reader to consume poisonous mushrooms was not subject to products liability laws. See Reutiman, *Supra* Note 101 for further analysis.

¹⁰⁴ Scott, *Supra* Note 102, 434-436.

¹⁰⁵ *Ibid.*, 461-462, referring to La Rossa v. Scientific Design Co 402 F.2d 937 (3d Cir. 1968) which stated: "[P]rofessional services do not ordinarily lend themselves to the doctrine of tort liability without fault because they lack the elements which gave rise to the doctrine."

¹⁰⁶ Retail Systems, Inc. v. CNA Insurance Companies 9 469 N.W.2d 735 (Minn. App. 1991) "The data on the tape was of permanent value and was integrated completely with the physical property of the tape. Like a motion picture, where the information and the celluloid medium are integrated, so too were the tape and data integrated at the moment the tape was lost."

¹⁰⁷ Saloomey v. Jeppesen & Co., 707 F.2d 671 (2d Cir. 1983); Halstead v. United States, 535 F. Supp. 782, 791 (D. Conn., 1982).

¹⁰⁸ Fluor Corp. v. Jeppesen & Co., 5 170 Cal. App. 3d 468, 216 Cal. Rptr. 68 (1985) (where the court referred to a previous holding that only innately dangerous items might be subject to products liability- and that therefore a sheet of paper might not. The California Courts of Appeals in Flour Corp. held that the potential danger posed by errors in the sheets of paper was sufficient to subject the object – aeronautical charts in that case- to products liability).

¹⁰⁹ To take "Waze" as an example (which will be further analyzed in the following chapter), Waze is embedded in the user's cell-phones, has 100 million active users,

It is therefore very likely, at least *prima facie*, that "thinking algorithms" too might find themselves classified as "products," even if their entire essence is information and even if their function replaces human services. In that sense, applying products liability rationales to distinguish when such algorithms should be governed by products liability makes sense.

B. On "Defects"

It was pointed out in Part II that as indicated by their titles manufacture defect and design defect both require the existence of a "defect." Thinking algorithms, however, are inherently expected to cause damage regardless of any defects. This is because sophisticated systems, in particular self-learning algorithms, rely on probability-based predictions,¹¹⁰ and probabilities by nature inevitably "get it wrong" some of the time. To take a concrete example: a medical algorithm is designed to diagnose patients and prescribe optimal treatment. Assume that the system is 100% certain that the patient has a type of disease that is cured without any intervention in 99% of the cases. But in the remaining 1% of the cases, the patient will die if not given treatment.¹¹¹ Given the probabilities of success, in 99 of 100

and, as elaborated on in *Supra* Note 137, is of clear potential danger. See Greg Sterling, *Waze Launches "Local" Ads Primarily Aimed at SMBs and Franchises*, Search-Engine-Land (Mar 28, 2018), available at: <https://searchengineland.com/waze-launches-local-ads-primarily-aimed-at-smbs-and-franchises-295285>).

¹¹⁰ A Netflix algorithm recommending movies does so based on a numeric prediction that we would like said choice, relying on our previous taste and an analysis of enormous databases of other consumers' preferences (Anthony Schneck, *The Subliminal Trick Netflix Uses to Get You to Watch Its Movies and Shows*, Thrillist Entertainment (Jan 8, 2018), available at: <https://www.thrillist.com/entertainment/nation/how-new-netflix-recommendation-algorithm-works>). A bail algorithm recommending whom to release and whom to deny bail, does so based on the probability that the suspect would break the law or escape if allowed to post bail. An application for choosing an optimal treatment for a patient too is based on the probability that the patient indeed has the medical condition diagnosed; and that she would react to the optimal treatment as most other patients would (Vigmesh Ramachandran, *Are Algorithms a Fair Way to Predict Who'll Skip Bail?* Futurity (Jun 5, 2017), available at: <https://www.futurity.org/bail-bias-algorithm-1450462-2/>).

¹¹¹ In real life, of course, the probabilities that are factored in are not one dimensional as in this example but are reflected in various stages of the decision-making process. Diagnosing the patient's medical condition is in itself often probability-based. Even if all signs indicate of cancer, for example, there is still a chance that the patient suffers from a different disease. (See, e.g., Lucy McNally, *Woman Endures Months of Unnecessary Chemo Treatment After Being Wrongly Diagnosed with Cancer*,

cases the algorithm would be right to choose not to intervene. The algorithm, however, naturally cannot "tell" in advance whether it is dealing with the 99 "ordinary" cases, or whether the patient before it is the 1 in a 100 exception. Choosing the optimal course of action will therefore be based on damage expectancy: which decision, if taken a large number of times, will lead to the best results? Assuming the magnitude of damage caused by no intervention is "0" for 99% of the cases and is "100" for the remaining 1%, and that the magnitude of damage due to unnecessary intervention is "10" (assuming it leads to significant side effects), the optimal decision will be not to intervene (as damage expectancy of intervening is $99 \times 10 = 990$, whereas damage expectancy of not intervening is $1 \times 100 = 100$).¹¹² Yet every once in a while our medical algorithm will undoubtedly encounter some "exceptional" patients as well, for whom its optimal choice of no intervention will be catastrophic.

Focusing on the damage caused due to a user being "on the bad side of the statistics" certainly does not mean that the system was defectively manufactured or defectively designed. On the contrary: the system has reached the decision we would want it to reach. It just so happens that whenever thinking algorithms reach decisions based on probabilities - which is exactly what they are designed to help humans with - inevitable damage will occur when the general rule is applied in cases that in hindsight turned out to be the exceptions.¹¹³ Does this mean that thinking algorithms should never be governed by products

ABC News (Nov 18, 2016), available at: <http://www.abc.net.au/news/2016-11-18/woman-given-unnecessary-chemo-treatment-bad-cancer-diagnosis-nsw/8036438>) Moreover, even if he does suffer from cancer, diagnosis of the specific type of cancer rather than another one is too probability-based. Even within the same types of cancer, different types of tumors exist, and determining between them might be a matter of probabilities.

¹¹² Ignoring the additional damage of "20" sustained by the 1% of patients who did need the intervention, since for them the other alternative is worse.

¹¹³ Note that because choices are made based on damage probabilities, it is also possible that the decisions chosen would in fact be damaging in the vast majority of cases, and that this would still be deemed the "right" one to make. If, for example, there is a 1% probability of catastrophic damage of "100" (such as death), and a 99% probability of an average damage of "1" (such as a minor scratch), a thinking algorithm would be correct in preferring to cause scratches 99% of the times (with damage expectancy of 99) rather than to cause death 1% of the times (with damage expectancy of "100"). In such cases the algorithm's choices would almost always be damaging, but in no way defective.

liability and that our analysis should have nothing to do with said legal framework? Not necessarily.

First, in addition to damages caused through no defect, thinking algorithms may certainly be responsible as well for defect-based damages, which do not stem from the user "being on the bad side of statistics."¹¹⁴ Secondly, "thinking algorithms" are not unique in their ability to cause damage in the absence of a defect. Traditional products too may be "defect-free" yet nevertheless cause damage.¹¹⁵ Indeed, to win a products liability lawsuit a plaintiff must prove - even when subject to strict liability theory - the existence of a defect,¹¹⁶ implying that certain damages are not caused by a defect.

Granted, one could argue that in thinking algorithms damage caused without the existence of a defect is *inherent*, unlike traditional products which do not inherently cause damage when no defect is involved. If so, products liability may indeed not be the most efficient framework applicable, especially given that products liability procedures are considered expensive and slow.¹¹⁷ While a different legal framework - whose underlying assumption is that no defect exists in the first place - might be more efficient, this does not render our current products liability regime irrelevant altogether.

Thinking algorithms, despite their nature as "information-based" and although they may frequently cause damage regardless of a defect, may thus nevertheless be governed by products liability. Now at last it is time to turn to the part where we examine how such algorithms reconcile with the rationales behind products liability laws, so as to draw up a list of specific parameters that would help us

¹¹⁴ If, for instance, the medical algorithm mentioned above had chosen to *intervene*, despite the lower chances of success and the higher damage expectancy, we would probably consider it a defect. By the same token, if the algorithm had reached the probabilities described above because it ignored critical information that was available, this would probably also be classified as a defect.

¹¹⁵ Car tires, for example, may explode after a certain time of usage and cause lethal damage. As long as the manufacturer provides appropriate warning as to their maintenance and frequency of replacement, any damage caused by such worn-out tires will not be attributed to a defect. See, e.g., *Carmichael v. Samyang Tires, Inc.*, 923 F. Supp. 1514 (S.D. Ala. 1996) pertaining to an accident resulting from tire failure ("*To maintain a claim under the AEMLD, a plaintiff cannot simply prove that an accident occurred and that he was injured; rather, "a defect in the product must be affirmatively shown"*").

¹¹⁶ *Ibid.*

¹¹⁷ See, e.g., Robert W. McGee, *Who Really Benefits from Liability Litigation?* Dumont Institute Policy Analysis No. 24. (1996), Available at SSRN: <https://ssrn.com/abstract=82596>. This is even more so in the context of algorithmic decision makers. See, e.g., Gurney, *Supra* Note 89, 262-264; David C. Vladeck, *Machines Without Principals: Liability Rules and Artificial Intelligence*, 89 Wash. L.Rev. 117, 137-141 (2014).

distinguish when products liability should apply and when it should not. To that end, let us first look at two sets of concrete examples of "traditional" versus "thinking" systems, which we will use throughout the remainder of the analysis.

PART IV: DISTINGUISHING PRODUCTS FROM THINKING ALGORITHMS

As discussed above, this article proposes an alternative approach for distinguishing traditional products from thinking algorithms which is not based, per se, on the system's level of autonomy. Rather, this part analyses how different features or characteristic of different decision-making systems – manifested in the four different OODA Loop stages– affect the achievement of the different rationales behind products liability. Practically, the more of the system's features reconcile with achieving products liability rationales, the more inclined we would be to classify them as traditional products. Systems whose features impede products liability rationales, however, should be classified as thinking algorithms that warrant different treatment. To identify and analyze the different features relevant to said "purpositive interpretation" analysis, let us look into two sets of concrete examples of decision-making systems.

A. *Examples of Products versus Thinking Algorithms*

One of the most notorious tragedies in the history of medical devices is that of Therac-25, a radiation therapy machine used to destroy cancerous tissues. Between 1985 and 1987, six patients in the United States and Canada were inadvertently given an overdose of radiation, resulting in three fatalities.¹¹⁸ Investigation revealed that the system had several "bugs" causing it accidently to release much higher dosages of radiation than prescribed by the machine's technician.¹¹⁹

¹¹⁸ Nancy G. Leveson and Clark S. Turner, *An Investigation of the Therac-25 Accidents*, Computer 18-41 (1993); Bev Littlewood and Lorenzo Strigini, *The Risks of Software* Scientific American 62-75 (1992); Johnathan Jacky, *Programmed for Disaster*, The Sciences 22-27 (1989).

¹¹⁹ For example, one of the software "bugs" that caused several accidents was that when the human operator inserting the treatment dosage made a change within a

Although the lawsuits filed in connection with the Therac-25 accidents were all settled before trial,¹²⁰ Therac-25 was considered one of the first cases to give rise to products liability claims in connection with medical devices.¹²¹

Therac-25 was useful for administering radiation in a precise and automatic manner; the example we shall use as its "sophisticated" counterpart is a machine also capable of taking and implementing professional decisions. Nowadays, "new generation" radiation machines mainly focus on improved precision of the radiation's distribution. Equipped with infrared cameras and robotic beds, radiation machines now make automatic minor adjustments in the positioning of the patient vis-à-vis the radiation beam throughout the radiation process, thus achieving more precise administering of the treatment.¹²² An additional feature that could be embedded in radiation machines, however, would also include dose calculation algorithms that would enable the machine to administer radiation beams but also to decide (or recommend) the optimal treatment plan for each patient based on his unique characteristics.¹²³ True, existing algorithms for calculating radiation dosage are in use today,¹²⁴ but the future system we take as an example makes use of "learning algorithms" that produce personalized dosage calculations based on the type of tumor involved, but also on myriad other parameters that the system itself deems relevant, after having "learned" from large databases of previous cases and itself having deciphered correlations between different parameters and improved outcomes.

time window of 8 seconds where it took the Therac 25 to set its magnets for operation, the change was not registered. Leveson and Turner, *Ibid.* 27-28.

¹²⁰ Sara E. Dyson, *Medical Device Software and Product Liability: an Overview (Part I)*, MedTechIntelligence (Sep 15, 2017), available at: https://www.medtechintelligence.com/feature_article/medical-device-software-products-liability-overview-part/2/; Hanna Bergman, *Medical Device Software: Who is Responsible When Something Goes Wrong?* OhioTiger (Oct 17, 2017), available at: <https://www.ohiotiger.com/medical-device-software-defects/>.

¹²¹ Dyson, *Ibid.*

¹²² Next generation state-of-the-art radiation therapy system at Netcare N1 City Hospital, Nercare, available at: <https://www.netcare.co.za/News-Hub/Articles/articleid/639/next-generation-state-of-the-art-radiation-therapy-system-at-netcare-n1-city-hospital>.

¹²³ Which consists of "the prescribed dose level for the tumor, the number of therapeutic beams, their angles of incidence, and a set of intensity amplitudes." U. Oelkfe and C. Scholz, *Dose Calculation Algorithms* in NEW TECHNOLOGIES IN RADIATION 187-196 (Schlegel et al. eds., 2006).

¹²⁴ *Ibid.*

To give some more background, algorithms' advancing self-learning abilities, coupled with enormous amounts of data, allow algorithms to learn from existing information and implement the conclusions in future sets of data. "Supervised learning" means a process where algorithmic training is more structured, in the sense that algorithms are fed with "right and wrong" answers pertaining to existing databases, so that they can develop a model for predicting the right answer for similar data sets that were not included in the training. A greater degree of freedom to come up with their own "conclusions" or recognize their own "patterns" is given to algorithms in the "unsupervised learning" process. Here algorithms are not "fed" any answers but are "free" to decipher patterns in the data that may indicate the right answer.¹²⁵ Indeed, in the field of radiation oncology concerns have been raised regarding the possible pace of developing machine-learning abilities, given the difficulty of collecting standardized data sets that the algorithms could train on.¹²⁶ Other projections, however, view radiation oncology as a very good candidate for personalized treatment based on self-learning algorithms.¹²⁷ Our first example of a "thinking algorithm" therefore focuses on a radiation machine that exhibits such projections.

Our second set of a traditional product versus a thinking algorithm takes us to the world of driving, more specifically the world of navigation. Global Positioning Systems (GPSs) were gradually put to civilian use at the end of the last century,¹²⁸ and are now commonly

¹²⁵ Avigdor Gal, *It's a Feature, not a Bug: On Learning Algorithms and What They Teach Us*, Roundtable on Algorithms and Collusion (Jun 21-23, 2017), available at: [https://one.oecd.org/document/DAF/COMP/WD\(2017\)50/en/pdf](https://one.oecd.org/document/DAF/COMP/WD(2017)50/en/pdf); Harry Surden, *Machine Learning and the Law*, 89 Wash. L. Rev. 87 (2014).

¹²⁶ Babette Ten Haken, *Radiation Oncology, AI and Machine Learning in Research*, (Jul 26, 2017), available at: <https://babettetenhaken.com/2017/07/26/radiation-oncology-ai-machine-learning/>

¹²⁷ Huan-Hsin Tseng et al., *The Role of Machine Learning in Knowledge-Based Response-Adapted Radiotherapy* in MACHINE LEARNING WITH RADIATION ONCOLOGY BIG DATA (Forthcoming), available at: <https://www.frontiersin.org/articles/10.3389/fonc.2018.00266/full>
<https://www.frontiersin.org/research-topics/6126/machine-learning-with-radiation-oncology-big-data>

¹²⁸ John E. Woodard, *Oops, My GPS Made Me Do It! GPS Manufacturer Liability Under a Strict Products Liability Paradigm When GPS Fails to Give Accurate Directions to GPS End-Users*, 34 U. DAYTON L. REV. 429 (2009).

used in different modes of transportation, including cars. By automatically determining a user's current location, the GPS uses its database of maps to calculate a route from said location to the end point as entered by the user.¹²⁹ In recent years numerous accidents have been reported to be the result of following GPS instructions: in addition to unfortunate encounters of vehicles with speeding trains,¹³⁰ GPS systems have also led drivers into a creek or to a cliff edge,¹³¹ instructed them to enter a road in the wrong direction or drive under a bridge too low for their vehicles to pass under safely;¹³² they have even guided them straight into a war zone.¹³³ Though the classification of GPS systems as "products" is not necessarily clear-cut, numerous courts have acknowledged the somewhat analogous object of aeronautical charts as "products" for the purpose of applying products liability.¹³⁴ A GPS is therefore the system we will use as an example of a "traditional product" for our analysis.

Our equivalent example of a more sophisticated system is Waze, a community-based navigation application designed for navigation per se, but also to "outsmart traffic."¹³⁵ Founded in 2006, the Israeli development makes use of real-time updates from its

¹²⁹ Ibid.

¹³⁰ Ibid.; Yaron Steinbuch and Hoa Nguyen, *Driver Blames GPS for Crash with Metro-North Train in Bedford Hills*, (Sep 30, 2008) available at: <https://www.ble-t.org/pr/news/headline.asp?id=23809>;

What if a GPS Causes an Atlanta Accident and Injuries? (Jan 19, 2017), available at: <https://www.atlantainjuryattorneys-blog.com/2017/01/gps-causes-atlanta-accident-injuries.html>

¹³¹ Edward J. Lake, *Defective GPS Systems and Product Liability for Accidents*, AVVO (Nov 21, 2012), available at: <https://www.avvo.com/legal-guides/ugc/defective-gps-systems-and-product-liability-for-accidents>; Mark Havnes, *Trusting GPS, Convoy of Visitors Get Monumentally Lost*, The Salt Lake Tribune (Aug 4, 2008).

¹³² Ibid.; *\$15M Products Liability Lawsuit Blames GPS Maker Garmin for 2013 Massachusetts Bus Crash*, Boston Car Accident Lawyer Blog (Jan 27, 2015), available at: https://www.bostoncaraccidentlawyerblog.com/2015/01/15m_products_liability_lawsuit.html

¹³³ Woodard, *Supra* Note 128; Yuval Azoulay, *American Tourist Stoned by Mob after Accidentally Entering Qalandiya*, Haaretz (Jun 26, 2008) available at: <http://www.haaretz.com/hasen/spages/996229.html>.

¹³⁴ See further discussion in Part III. See also *Turner v. iSecureTrac Corp.*, No. 03CA702004 WL 944386 (Ohio App. 5th Dist. Apr. 28, 2004), where the Fifth District Court of Appeals court in Ohio did not question whether a GPS ankle device was a "product" under products liability laws, as well as *Cruz v. Talmadge*, 244 F. Supp. 3d 231, 2017 US Dist. where the Massachusetts United States District Court remanded to the superior court department a case involving products liability claims due to GPS navigation orders that caused a bus crash into a bridge whose height was below the height restriction for buses.

¹³⁵ <https://www.waze.com/>

community of drivers – automatic ones sent by the system itself reporting its location, and traffic submissions actively sent by users – in order to create an ever-updating mapping of the roads system.¹³⁶ Thus, Waze's algorithms are able to calculate the shortest way from point A to B, but also the quickest way at any given moment. As successful as it is, Waze too has been blamed for leading drivers into dangerous situations that might result in harm.¹³⁷ As we shall soon see, various features of Waze render it different from traditional GPSs, when the system's compatibility with products liability laws is analyzed.

B. *Which Products Liability Rationales Apply to Which Algorithms*

For an orderly analysis of the compatibility of the various rationales of products liability laws with both sets of examples, the article addresses each rationale separately and examines how the thinking algorithm's features affect it as compared with the traditional product.

1. Promoting Safety

As discussed above, a central rationale behind the framework of products liability is to encourage manufacturers to better their products' safety. Whether or not products liability indeed achieves said end has been the subject of heated debates, but for the sake of discussion we shall assume that in general it does.¹³⁸ But how is

¹³⁶ <https://web.archive.org/web/20130708113343/http://world.waze.com/faq/>

¹³⁷ *Israeli family accidentally enters Ramallah*, Arutz Sheva (Feb 24, 2018), available at: <https://www.israelnationalnews.com/News/News.aspx/242361>; Chana Roberts, *Three Israelis accidentally enter Ramallah*, Arutz Sheva (Dec 10, 2016), available at: <https://www.israelnationalnews.com/News/News.aspx/221510>

¹³⁸ According to opponents of the products liability regime, empirical evidence indicates that products liability has not been shown to significantly increase the safety level of products. This is so, the argument goes, because in most cases market powers as well as regulatory measures are themselves responsible for increased safety, such that the incremental addition in safety attributed to products liability litigation is non-significant. Polinsky and Shavell, *Supra* Note 82.

Either way, both proponents and opponents of the products liability regime agree that when the public and regulators' level of awareness of the danger or damages caused by the product is low (for instance, because the product is not widely sold, because the media coverage of the damage is not significant, because the damages

"promotion of safety" affected when the "products" at hand are sophisticated self-learning systems?

Generally speaking, manufacturers of such systems might find themselves in a place where enhancing safety is very difficult or expensive, or will render the system inefficient. This point may be explained through the example of a futuristic system of a "robo-doctor" able fully to replace a human physician at all stages of decision-making.¹³⁹

First, manufacturers of traditional products have a finite number of parameters to consider when preparing for different scenarios and minimizing risks associated with the actions of their machines.¹⁴⁰ By contrast, robo-doctor manufacturers will have an infinite number of scenarios against which they must try to take precautionary measures, based on innumerable parameters: the patient's general medical condition (blood type, vital signs, height, weight, etc.), past medical condition (previous lab results, previous diagnoses, previous success or failure of past treatments, etc.), as well as current medical condition (e.g., for cancer diagnosis: type of cancer, its size, its location, its stage, which receptors it has, etc.); various external parameters (are any epidemics indicated in that region? Do current weather conditions affect the likelihood of a certain diagnosis

are difficult to trace back to the product or are only discovered after a long time, etc.), regulation and market forces will not suffice to promote optimal safety. *Ibid.*; Goldberg and Zipursky, *Supra* Note 69. General technological and social trends may result in damage being more easily and quickly discovered, and gaining more public exposure. To take the damage caused by the Therac-25 as an example, with time improved safety mechanisms would likely have allowed the machine to check the dosage it administered in practice, and thus immediately discover discrepancies. In addition, nowadays reports on over-dosage caused by a radiation machine would probably rapidly 'go viral' due to globalization and the social media effects (as opposed to the Therac-25 accidents that occurred in the 1980s of which the public was not aware for many months). Indeed, systems causing damage nowadays are likely to attract more public and regulators' attention than did systems that caused harm decades ago. In this sense, products liability in general perhaps is less critical today for increasing safety than it was in the past. But said difference does not stem from the difference in function or abilities of traditional products versus sophisticated ones but from general trends. As such, they are beyond the scope of our current discussion.

¹³⁹ Alex Woodie, *The Robo-Doctor Is [In]*, Datanami (Au 30, 2017), available at: <https://www.datanami.com/2017/08/30/the-robo-doctor-is-in/>; Michael MacRae, *The Robo-Doctor Will See You Now*, ASME (May 2012), available at <https://www.asme.org/engineering-topics/articles/robotics/robo-doctor-will-see-you-now>.

¹⁴⁰ To take a simple example of a coffee machine: verifying that the temperature of the liquid produced by the machine is not too hot, verifying that even if operated by a child the machine could not cause electrocution, etc.

or the chances of success of a certain treatment? Should a certain blossom in the air be taken into account when assessing the reason for a symptom?); practical parameters (are qualified staff available immediately to execute a certain medical choice? What is best practical choice in cases of understaff, or of shortage of specific medications in the hospital's stock?); ethical parameters (what does the patient truly want? How does it reconcile with ethical standards as well as the relevant legislation?).

Secondly, in addition to all the foregoing parameters and scenarios, manufacturers have to deal with information fed into the system by external systems (especially considering the Internet of Things revolution, and the prospect that in the future algorithms will rely on parameters fed to them by other machines—such as a robo-doctor factoring blood pressure and heartbeat rates broadcast to it by a separate medical device, and perhaps having to assess the probability that said device has provided erroneous readings).¹⁴¹

Thirdly, medicine, like other professional fields, changes constantly. Do we expect manufacturers to have their products updated daily (or even hourly) with every new study, while immediately incorporating that study's findings into the robo-doctor's decision-making process? Who will decide which studies should be updated and which are not convincing enough, or are less relevant to the specific population treated by that robo-doctor? Having to account for such dynamic developments, the manufacturers' task in minimizing risk of error is undoubtedly far more difficult than in the case of "static" fields where products need not be subject to frequent updates.¹⁴²

Fourthly, this also leads to another unique challenge for manufacturers of thinking algorithms: medicine, like law and many other complex fields where judgment and discretion are significant, is not black and white. Different experts have different opinions and recommend different solutions under identical circumstances. How

¹⁴¹ See, e.g., Adam Thierer, *The Internet of Things and Wearable Technology: Addressing Privacy and Security Concerns without Derailing Innovation*, 21 RICH. J.L. & TECH. 6, 4-17 (2015).

¹⁴² Complex as an autopilot system might be to design, for example, its manufacturers are not likely to encounter daily updates in the field of aviation (or meteorology, atmospheric science, etc.) that would require them to decide whether the new information is relevant to the design of the system and if so how the system should take the updates into account.

can manufacturers be expected to minimize the risk of certain scenarios by choosing a specific course of action, when the choice is not necessarily obvious? (In any event, if the matter is ever litigated the manufacturer could always be blamed for not choosing a different alternative that some experts preferred.) Lastly, even if raising the safety level despite the countless scenarios, manufacturers have to prepare for would be possible (and not prohibitively costly). Even disregarding the added complexity of the "no right or wrong" answers, manufacturers of thinking algorithms will still have a very difficult time improving the rates of correct decisions of their systems from "high" to "very high" if they wish to keep the system efficient and user-friendly. For example, if the robo-doctor has to include in its decision-making process each and every possible medical condition, including extremely rare ones or implausible ones (in order to reduce the chances of misdiagnosis, even for the most unusual cases), it might not be able to be put to practical use. The extra time required for the system's information gathering and analysis process might make the process excessively long; make the system fail to respond in real-time; make patients refuse to be anemized if the process is so slow; or make hospitals refuse to buy more robo-doctors to compensate for each unit being capable of treating only few patients.¹⁴³

In sum, the factors affecting how much a manufacturer of a sophisticated system could indeed increase its safety are these: the size of the matrix of parameters the algorithm must consider before making a decision;¹⁴⁴ the dynamic nature of the relevant professional knowledge; the lack of clear "right choices"; and the extent of tradeoff between safety and efficiency. But these parameters are all very general, and deduced from an extreme example of a robo-doctor - a

¹⁴³ On the inherent trade-off between safety and efficiency in the field of robotics see e.g., Christiana Braz et al., *Designing a Trade-Off Between Usability And Security: A Metrics Based Model*, in HUMAN-COMPUTER INTERACTION – INTERACT (Baranauskas et al., eds, Springer 2007). For a discussion on said trade-off in the pharmaceutical field see, e.g., Tomas J. Philipson and Eric Sun, *If The Food and Drug Administration Safe and Effective?* 22 *Journal of Economic Perspectives* 85-102 (2008).

¹⁴⁴ An added complication that increases the magnitude of scenarios a manufacturer must prepare for is when parameters are not binary but can be anywhere on a spectrum; This is even more so when the parameters are not deterministic but in themselves reflect probabilities. The parameter of a patient's past medical history of a stage 3 carcinoma, for example, does not reflect one hundred percent certainty that indeed this was the exact medical condition the patient suffered from. Rather, it reflects a certain probability that indeed the diagnosis was correct.

machine that wholly replaces one of the most complicated human professions. To try and concretize the parameters that affect manufacturers' ability to increase safety, hence to meet the first rationale of products liability, let us now analyze our set of two more delineated examples of radiation and navigation systems. The analysis centers on the level of *foreseeability* of the product's "actions," as well as how far the manufacturer is able to *control* said actions, and do so *efficiently* (assuming that the less a manufacturer is able to foresee and control the choices of its system, the less will it be able to raise the system's safety levels). To examine this, we shall review separately the system's operation in each of the four OODA Loop stages.

i. The OODA Loop - General

We might not have thought about it this way, but even primitive automatic systems may be responsible for all four stages of the OODA Loop (provided they yield some sort of physical change). Even a simple drill, for example, "acquires information" about the fact that it has been turned on, and that the human operator has pressed the right button for the drill to start. It also analyzes the information and makes a decision based on it, in the primitive sense of considering that the "start button" was pushed, along with the fact that no safety lock mechanism was initiated, and then it "decides" to "go ahead" and drill, the drilling being the final OODA Loop stage of execution. But in such examples all three first stages of the process are injectively dictated to the system by its human operator, leaving full foreseeability and full control in the manufacturer's hand (unless a "bug" occurs).

"Thinking algorithms" however replace humans precisely because not only can they outperform them in the final physical execution stage, they can also automatically access and collect vast amounts of information from various sources, of a magnitude that the human brain could not read in decades; can analyze these enormous amounts of information that are beyond a human's grasp; and can make complex decisions based on probabilities that a human cannot even weigh.

The Therac-25 machine, for instance, only acquired information that was 100% dictated to it by the operator (having turned the machine on, chosen a specific treatment mode, inserted the desired parameters for treatment, etc.). It "analyzed" whether it was in a

position to start operating, based on pre-programming that injectively ordered it to do so when all conditions for beginning an operation mode were met, and then "blindly" "decided" to go ahead and administer treatment - again because it was programmed to do exactly that once all the required conditions were met. A new generation radiation machine, however, would have more "freedom" or "independence" to conduct the first three stages of the OODA Loop, such that the manufacturer would no longer have full foreseeability or control over them. With respect to information acquisition, for example, it is possible that the system itself will decide which sources of data to harvest to improve its success rates (be they medical publications the manufacturer might not even be aware of, or random statistics on such matters as global warming that the system might suddenly find are correlated to radiation success rates). With regard to the analysis stage too, the radiation machine itself might determine how much weight to attach to each piece of information it has collected, again based on correlations it itself has discovered through its self-learning process. Facing several courses of action ranked with different probabilities of success and different expected damage, the machine itself might decide which alternative to choose, or otherwise decide that its confidence in the preferred option was not high enough. Then it will rather call a human for further instructions than decide to execute.¹⁴⁵ Naturally, the more OODA Loop stages a system is capable of performing in a manner not fully dictated by the manufacturer, the less foreseeability the manufacturer has as to the final outcome of the process, and the less control it obtains over it.

Therefore, and before delving into each of these stages separately, our first parameter indicating that the system is a "thinking algorithm" that is less compatible with the first rationale of products liability (given that reduced levels of control and foreseeability would render it more difficult to improve safety, as explained above) is the number of OODA Loop stages the systems performs in a manner not fully dictated by humans.

ii. *"Observe" ("Information Acquisition")*

As explained above, the information acquisition stage of Therac-25 was fully dictated by humans in the loop. A new generation radiation machine, depending on its specifications, might raise several separate aspects of lack of foreseeability and control by the manufacturer at the information acquisition stage. Naturally, when the system trains on closed sets of databases that the manufacturer has "fed" it, and later continues to collect information from data provided to it exclusively

¹⁴⁵ See stage 5 of Sheridan's spectrum, *Supra* Note 51.

by the manufacturer, the manufacturer maintains foreseeability and control over the information acquisition stage. However, better results might materialize if the system is free to decide on its own to add additional sources of information (for instance, additional medical journals, readers blogs on radiation, etc.) as well as more types of information the system is "interested in" (for instance, when reviewing a patient's medical record, collecting information on less trivial types of information such as the day of the week the patient was released from the hospital, etc.). In such cases lack of foreseeability by the manufacturer is threefold: first, it cannot anticipate the *type of parameters* the system will choose to collect information on. Secondly, it cannot anticipate the *sources of information* the machine will harvest. And thirdly, it cannot anticipate the *specific content* of the data collected, be it content related to external databases such as medical journals (where the more dynamic and routinely updated the database is, the less likely a manufacturer is to have foreseeability over its content) or the values of the specific medical parameters measured for each patient.¹⁴⁶

Foreseeability and control at the information acquisition stage, therefore, are reduced if the system can decide on its own what type of information to seek or which sources of information to cover, and the more dynamic its information sources are.

To move to our second set of examples, from the navigation world, here both the traditional GPS and its "sophisticated" Waze counterpart are engaged in the information acquisition stage. In the case of the traditional GPS the manufacturer has full foreseeability and control over the data acquired, since the system is based only on maps uploaded by the manufacturer itself (as a one-off occurrence when the system is released to the market; through updates that the manufacturer "sends" to the system; or the user downloads from the manufacturer's website).

¹⁴⁶ The manufacturer, for example, cannot foresee that a certain patient will be O+ blood type rather than A+ (particularly if said blood type is not much more common than the others). It can anticipate, however, that the patient will be one of the A+/A-/B+/B-/AB+/AB-/O+/O- blood types, and not, say, a whole new blood type M+. In addition to limited foreseeability, control over said content is also problematic, in the sense that the manufacturer does not dictate the specific values of the patient but also in the sense that the system might be fed erroneous information beyond the manufacturer's control (potentially by other systems connected to it via the internet of things).

The manufacturer, however, does lack control over information that has *changed* from the time the maps were updated until the actual usage of the system (roads closed since the last update; regions that have become hostile; bridges built with a maximal height cap; etc.). Likewise, manufacturers will probably have limited foreseeability with respect to such changes (major changes may well be planned and advertised in advance and thus are foreseeable, but small or temporary changes might be executed without the manufacturer learning about them in advance).

How do the foreseeability and control elements change when the system is not based on maps uploaded by the manufacturer, but on information sent by the users themselves in real time?

Naturally, the system's being based on real-time reports sent by millions of strangers limits the amount of foreseeability of said reports. (The manufacturer can assume, for instance, that reports are unlikely to come from vehicles driving in mid-ocean; but in general which new routes users will report driving on, etc., cannot be anticipated.) Likewise, the manufacturer has no control over the content of information received (later, at the information analysis stage, it may decide to disregard certain types of information that it suspects are unreliable. But the content of information sent and acquired is within the control of the system's users, rather than its manufacturer's).

In fact, not only does the manufacturer have limited foreseeability as to the content of information the system relies on, and not much control over it, but the communal nature of the information acquisition stage makes the system vulnerable to *intentional false reports*. (For instance, users wish to "game" the system to divert other drivers from their own route and reduce traffic;¹⁴⁷ psychopaths may attempt to "swamp" the system with false reports of a new road built over a lake, only to have the system plunge innocent drivers into it.) The manufacturer can take precautions to try to discard false information (accepting information only from users who previously have provided information that proved true; accepting information only when a set minimal number of users support it; accepting certain types of potentially dangerous pieces of information only after a Waze representative has physically visited the location to verify it; etc.). Precautions such as these are relevant in the context of the next OODA

¹⁴⁷ Allisa Walker, *Is It Really Possible To Trick Waze To Keep Traffic Off Your Street?*, Gizmodo (Nov 18, 2014), available at: <https://gizmodo.com/is-it-really-possible-to-trick-waze-to-keep-traffic-off-1660273215>. See also Nicholas Tufell, *Students hack Waze, send in army of traffic bots*, Wired (Mar 25, 2014) available at: <https://www.wired.co.uk/article/waze-hacked-fake-traffic-jam>.

Loop stage of information analysis, rather than information gathering where the system depends on input from users.

On the other hand, compared with traditional GPS, Waze minimizes the lack of foreseeability aspect associated with recent changes that occurred after the system was released, or at least shortens its duration. This is because the system learns of such changes in real time, and adds the missing information as soon as a sufficient number of reports is received.

So in general, and as may be intuitively expected, a system based on information received from its users is characterized by less foreseeability and control than in the first OODA Loop stage. On the other hand, a system based on information received in real time (rather than on information collected once, or with long spells between collections) is likely to suffer less from lack of foreseeability with respect to updates and new information.

iii. "Orient" ("Information Analysis")

As with the stage of information acquisition, Therac-25 was not involved in information analysis: the analysis was done externally, by humans. The new generation radiation system, on the other hand, has a significant role in the analysis stage, and as a result lowers the manufacturer's level of foreseeability and control over said stage. In more detail, the system's mere computational abilities, which enable it to weigh up the various complex parameters collected in the previous stage, do not impair the manufacturer's foreseeability or control. Even if the manufacturer cannot perform the computation tasks itself, the information analysis stage is perfectly foreseeable and controllable: the manufacturer is the factor that decides how much weight the system should give each parameter, and which collected pieces of data to disregard altogether. The manufacturer can decide, for instance, that a patient's age should only be considered when the patient is very young or old, or decide that much weight should be given to the fact that a patient has a history of prior tumors and the system will conduct its computational analysis accordingly. New generation machines, however, are likely to be more than sophisticated calculators. Rather, the whole uniqueness of advanced systems is that they can learn for themselves, at much better success rates than humans, how much weight to attach to each parameter based on prior experience. In such cases it is the machine that will decide how far to consider a patient's

medical history, her susceptibility to allergies, or a new controversial study published in a medical journal. The manufacturer's ability to increase control and foreseeability by setting boundaries is limited in the information analysis stage because it can do so mainly sporadically, and with respect to the weight to be given to certain parameters that the manufacturer knows in advance will be part of the system's analysis. Here too, the more dynamic the database the system draws information from, the less able a manufacture is to pre-instruct in a broad manner that certain weights be given to certain parameters.

On the other hand, the navigation sets of examples seem less revealing in the context of the second OODA Loop stage, as the differences between the traditional GPS and Waze are less evident.

In general, the manufacturer of a traditional GPS has ample foreseeability and control over the analysis stage of the system. The manufacturer dictates in advance how much weight be given to the different parameters (for instance, if a map shows that a road is closed, the manufacturer will likely instruct the system to attach 100% weight to said piece of information and disregard the existence of this road when calculating a route). Users themselves may also influence the analysis stage of the GPS (for instance, by instructing it to avoid toll roads, or to present the shortest route as opposed to the fastest route, etc.). But it is the manufacturer that designed these choices that users may make, hence these choices are foreseeable and controllable. As mentioned, traditional GPSs lack the information that has changed since the last update of the system; but precisely because such information is not part of the data that the system analyzes there is no sense in foreseeing or controlling *how* the system would analyze it because it simply will not.

With Waze, the real-time character of the updates might cause the system to attach different weights to similar reports, based on the volume of the reports (for example, the system is likely to give much weight to 300 reports of a traffic jam on a country road usually not taken by many cars, and only little weight to two reports of a traffic jam at a very busy intersection). The decision on how the weight ought to change based on the number of reports and other parameters, however, is predetermined by the manufacturer and in that sense is foreseeable and controllable. In addition, the manufacturer is free to decide to attach zero weight (in other words fully disregard) reports it deems "suspicious" as false.

Moreover, even in respect of real-time updates that will add new unknown information and will "force" the system to decide how much weight to attach to such pieces of data, this too can be predetermined. (For instance, if the new information indicates any sort

of danger, the system will attach maximal weight to it and verify that no driver is directed to the danger zone until it is proven safe; if the new information indicates an unfamiliar new road, the system will not direct users there until a certain volume of reports has been received; etc.)

The "wisdom of the crowd" and the "real-time" reports nature of Waze, therefore, do not seem highly significant as regards the foreseeability or controllability of the information analysis stage.

iv. *"Decide" ("Decision Selection")*

Having attached different weights to the myriad parameters collected through the information acquisition process, and having analyzed it, a system also responsible for the third OODA Loop stage now has to make a decision based on said analysis. This stage involves more than might meet the eye, at least for thinking algorithms. Unlike a coffee machine, for instance, whose decision-making process is generally based on two deterministic options, "pour coffee" or "don't pour coffee," a thinking algorithm might face numerous alternatives, each based on probabilities and each accompanied by a certain level of confidence that indeed said alternative entails said probabilities. After analyzing the information it has acquired, a new-generation radiation machine, for example, might come up with dozens of potential treatment dosages, each entailing different probabilities of success and of expected damage. The algorithm's ranking may, for instance, include an option whose success rates are 90%, entailing damages of "3000" for the 10% of failure; an option successful 80% of the times, entailing damages of "1000" in the 20% cases of failure; and a long list of alternative dosages of different success rates and expected damage. Not only that. Like a human physician, the algorithm cannot be 100% certain that said alternatives *indeed* reflect the probabilities and expected damages the algorithm assumes they do. The algorithm can, for example, determine that it is 95% confident that the first option indeed has a 90% probability of success and a potential of causing damages of "3000," and is only 70% confident that the second option indeed reflects the rates indicated. Our decision selection stage, therefore, involves some "tough" questions for an algorithm. First, does it give more weight to potential success rates or to potential damages? (Under the first alternative success rates are 90% and expected damages are $10\% \times 3000 = 300$. Under the second alternative success rates are only 80%, but on the other hand expected damages

are lower: $20\% \times 1000 = 200$. Which is preferable?) Secondly, how does confidence level affect the choice among the various options? Should the algorithm defer to a human whenever the rate of confidence of its preferred alternative is lower than a certain threshold? (Will the algorithm itself decide that it has to consult with a human, or will the threshold be predetermined by the manufacturer?)

A manufacturer's level of foreseeability and control naturally depends on how, or more precisely who, gets to decide these questions. With respect to the radiation machine the manufacturer may decide, for example, to adopt a more careful approach, where a machine is not free to decide on an alternative whose expected damages are more than a negligible percentage, and that whenever its level of confidence is below a very high threshold it must "step back" and let a human decide. Such an approach, however, will naturally be at the expense of efficiency (as in many cases the machine will not be able automatically to complete the process but will have to wait for a human to arrive and make her decision). Naturally, if the situation calls for instant treatment (for instance in a system used in a trauma unit) having to wait for a person might cost lives. Also, if from the outset the system's success rates are higher than the human counterpart's, from a utilitarian point of view we would prefer the machine to make such decisions, not the humans involved¹⁴⁸ (leading, again, to greatly reduced extents of foreseeability and control by the manufacturer).

The more the system's response time is critical, and the wider the gap between a human's and a machine's success rates (in favor of the machine), the more likely will the manufacturer be forced to forego foreseeability and control at the decision-making stage, and "free" the system to make its own choices. In other words, while the urgency required and the system's superiority over people do not in themselves render the system less foreseeable or controllable, such result it probable given efficiency considerations (as will also be discussed in the part addressing chilling effects).

Unlike a traditional GPS, which decides based on deterministic parameters (such as distance and existence of available roads), Waze makes its decisions based on probabilistic parameters that change constantly.¹⁴⁹ Being a "real-time service," Waze cannot wait for a human in the loop to assist it whenever its level of confidence is not sufficiently high, which again makes the system less foreseeable and

¹⁴⁸ Millar and Kerr, *Supra* Note 34, Lieblich and Benvenisti, *Supra* Note 58.

¹⁴⁹ Such as the degree of traffic expected at a certain location at a certain point in time when the user is expected to cross said location.

controllable. But unlike a radiation machine, Waze is not intended to save lives therefore enjoys the "luxury" of hedging risks. If the system has indications that a certain road is not safe (is under construction, leads to a danger zone, etc.) its manufacturer may theoretically decide in advance to eliminate such an option entirely until it is proven safe (at the expense of directing users to less efficient routes). Life-saving or medical algorithms, however, do not generally have this option, otherwise any choice they made would entail certain risk. Foreseeability and control (hence compatibility with the rationale of promoting safety) are thus reduced in such systems.

v. *"Act" ("Action Implementation")*

Neither of our radiation and navigation systems, whether the traditional or the "thinking" type, has ever replaced humans in the final OODA Loop stage of action implementation (GPS/Waze because they do not have an element of execution to begin with; Therac-25 and the new radiation system because the projection of the beams were never a process that a human performed). To touch on the stage of action implementation, let us therefore think of a robo-surgeon (for instance, the da Vinci system),¹⁵⁰ and focus solely on its execution rather than the decision-making process of how to execute. With respect to said specific stage, there is no "learning" element. Just like a traditional product, therefore, the system is foreseeable and controllable as long as it does not encounter "bugs."

vi. *Interim Summary & Measurability of Success Rates*

An analysis of the four stages of the OODA Loop and the levels of foreseeability and control associated with each in respect of our two examples of learning algorithms, revealed that the following parameters tend to reduce the system's compatibility with the rationale of encouraging manufacturers to promote safety: the greater number of the OODA Loop stages the system is responsible for; a system's freedom to decide which sources to draw data from; a system's freedom to decide which parameters to consider; the dynamic nature of the information sources relevant in the field of the system; use of "the wisdom of the crowd" to obtain information; the urgency of the process the system is involved in (or the requirement of real-time

¹⁵⁰ <http://www.davincisurgery.com>

actions); whether the system is life-saving (such that reduction of efficiency to increase foreseeability is very problematic); whether the system's success rates are already higher than those of a human equivalent (which again would render sacrificing efficiency to promote foreseeability more problematic) whether the system considers real-time updates by its users (interestingly, this characteristic may work both ways: on the one hand its content may "surprise" the manufacturer but on the other hand it may decrease unforeseeability associated with data that have changed without the manufacturer's knowledge).

Naturally, this is a rather initial list of parameters, and additional examples might yield additional ones. But it does give us a notion of the type of systems that should be classified as "thinking algorithms," in the sense that improvement of their safety by the manufacturer is more difficult or problematic to accomplish.

In addition to these parameters which impede improvement, another general parameter to take into account is how easily improvement is *measured*. The success rates of certain systems are fully measurable (annual number of car accidents autonomous vehicles are involved in per miles traveled; number of fatalities per miles traveled, etc.), but other systems' success rates may well measure the rates of "false negative" failures but not the "false positive" ones, or vice versa (failure rates of a bail algorithm which approved bail for defendants who later broke the bail rules are easily measurable. The rates of defendants who were denied bail, even though they would have not broken its rules if given the opportunity, are not). It may be difficult or almost impossible to measure success rates of other systems (for example, "success rates" of Waze, of algorithmic company directors, or of algorithmic priests). Measurability of success or failure does not directly contribute to manufacturers' ability to improve their systems' safety, yet it does give them a clear picture of the system's performance, and may encourage them to try alternative designs until better results are measured (even if manufacturers cannot explain why such designs work better than the previous ones, in light of lack of foreseeability and control difficulties).¹⁵¹ The extent to which the system's failure rates are

¹⁵¹ When Facebook decided to change its news feed algorithm for example, the team admittedly did not know how the change would affect the algorithm and its choices. Rather, to implement a change, many trial and errors experiments are required. Will Oremus, *Who Controls Your Facebook Feed*, Slate (Jan 3, 2016), available at

measurable, then, may assist manufacturers improve their systems' safety level, therefore indicates compatibility with the first rationale of products liability law.¹⁵²

2. Avoiding a Chilling Effect

It has been argued that in certain industries, the application of products liability not only failed to achieve an increase in safety (or in victim's compensation, which will be discussed next), but led to higher production prices resulting in a suboptimal level of manufacturing or use of beneficial technologies.¹⁵³ Although the avoidance of a chilling effect and the promotion of an efficient level of use are not one of the main rationales behind products liability but are rather competing interests, the shaping of the products liability framework was significantly affected by said interests.¹⁵⁴ Our analysis will therefore include said interest and examine how different parameters of sophisticated systems affect concern over a chilling effect.

In general, the literature addresses concerns that products liability will have a chilling effect on sophisticated systems. According to said concerns, the volume of products liability claims associated with such systems is expected to be high, given the increasing number of cases where damages are not caused by a human

http://www.slate.com/articles/technology/cover_story/2016/01/how_facebook_s_news_feed_algorithm_works.html.

¹⁵² For some systems, exact measurement of failure or success rates might be possible, but only after the passage of time. Mortality and morbidity rates associated with a new generation radiation machine, for instance, will only be revealed years after its release. During the time when statistics are unavailable, the "measurability" parameter will incline to the "non-compatible with products liability framework" but as soon as it becomes available the system can be treated as more inclined to products liability compatibility.

¹⁵³ According to Polinsky and Shavell, for instance, products liability led to an astronomic increase in the prices of the DPT vaccine, leading to significant under-vaccination. Similarly, litigation costs associated with products liability claims in aviation led to a dramatic increase in sales prices and to suspension of production by leading manufacturers. Polinsky and Shavell, *Supra* Note 82.

¹⁵⁴ Promoting an efficient level of usage, however, was discussed as a sub-rationale in the context of deciding between applying strict products liability versus a negligence-based one. In general, concern over a chilling effect swung the pendulum of the products liability framework in the latter direction. See discussion in Part II.

tortfeasor but by a system that has replaced her,¹⁵⁵ and since in general the introduction of new technologies has traditionally resulted in an increase of tort claims associated with it.¹⁵⁶ In addition, as a result of globalization and social media, manufacturers will probably receive more post-sale reports regarding the actual use of their products than in the past, advising them of additional potential risks not anticipated in advance. Such additional reports will impose warning requirements on the manufacturers which, when not followed, will be used against them in "failure to warn" products liability claims.¹⁵⁷

Naturally the less foreseeable and controllable a system is, the greater is the fear of a detrimental effect on technology. First, and as discussed above, lack of foreseeability and controllability (divided into its sub-characteristics) render it more difficult for a manufacturer to improve safety, at least for a subset of the innumerable scenarios possible. In such cases products liability will not necessarily contribute much to safety, but will likely result in higher production costs which will translate into reduced manufacturing of certain systems due to less demand. Moreover, lack of foreseeability is likely to render liability costs less predictable, and in turn again delay development or result in high costs.¹⁵⁸ Secondly, also as discussed above, naturally the better result a system has compared with a human equivalent, and the speedier its response, the more decisions will we be likely to entrust it with, leaving humans "outside the loop."¹⁵⁹ So for systems whose results show superiority over humans, fear of a chilling effect (in the form of decreased development, decreased

¹⁵⁵ Colonna, *Supra* Note 8.

¹⁵⁶ *Ibid.* (focusing on the biotechnological industry as an example of a "severe barrier for innovation" caused by products liability).

¹⁵⁷ See RESTATEMENT (THIRD) OF TORTS, § 10(b) (1998), that imposes post-sale warning requirements when the seller knew or should have known of a substantial risk of harm. Alternatively, it has been argued that the volume of products liability claims will not increase with the introduction of sophisticated decision-making systems but in fact might decrease, given the improved safety rates expected of these products, which will reduce the general occurrences of damage. Villasenor, *Supra* Note 19.

¹⁵⁸ Smith, *Supra* Note 19. Granted, for systems with a potential to cause damage that may be used by the user repeatedly, such as autonomous cars, we would want their prices to reflect a certain amount of risk for damage, in order for the user to internalize expected damages costs and use the system efficiently (for instance, not send its vehicle for rides around the neighborhood just as a joyride). Similarly and as discussed in Part II, one of the goals of the products liability framework is to have the manufacturer internalize the costs associated with the damages it might cause. Polinsky and Shavell, *Supra* Note 82. But the more unpredicted these costs are, and internalization of them is less likely to reflect the actual value of such costs.

¹⁵⁹ Abbott, *Supra* Note 27. Millar and Kerr, *Supra* Note 34.

demand due to high price, or reduced efficiency resulting from "clumsy" safety measures that render the system too slow or non-user-friendly) is of greater concern. To give a concrete example: to decrease products liability litigation risks, the manufacturer of the new generation radiation system may design the system to be more foreseeable. For instance, it will "feed" it with the information it learns from and not allow it to search independently for relevant sources of information; or it will heavily intervene and dictate the weight the system will attach to different parameters, not leaving the system discretion to decide based on its past experience. By doing so, manufacturers might increase foreseeability and control, but will not necessarily improve the system's outcomes, compared with older versions of the system or with humans. Assuming society wishes to promote the development of such a better technology (which is perhaps conceptually easier once the system already outperforms humans and is thus expected to save lives or reduce injuries), products liability might be over-burdensome on certain technologies.

3. Victims' Compensation

Generally, products liability has been accused of not being an optimal regime with respect to compensating victims for their damages, given the high costs associated with products liability litigation that render many damage cases unactionable¹⁶⁰ and lead to significantly reduced compensation for the victim in cases that are filed due to legal fees.¹⁶¹ How is victims' compensation affected by the characteristics of the damaging product? Unlike the case with the first products liability rationale of promoting safety, the particular abilities of each algorithm will likely not play a crucial part in affecting victims' ability for redress, although some more general characteristics will.

First, in many cases insurance of various types may render products liability redundant, as damages claims of insured victims might be covered in full by insurance, which eliminates the need for

¹⁶⁰ Given that the expected litigation costs exceed the expected reward for the damage suffered by the product, for instance, in car accidents where no significant physical injury was caused. Hubbard, *Supra* Note 86, 1826-1827; Gurney, *Supra* Note 89, 265-266.

¹⁶¹ Polinsky and Shavell, *Supra* Note 82 (reviewing empirical studies that found plaintiffs receive only 0.37–0.6 % of the amount paid by defendants).

litigation.¹⁶² Indeed, in the US certain types of insurance, including life insurance, health insurance, disability insurance, property insurance and car insurance, are prevalent.¹⁶³ A significant amount of damage cases caused by products are nevertheless uninsured, whether because not all victims have insurance or because the insurance does not cover the full amount of the damages suffered.¹⁶⁴ From the standpoint of assuring victims' compensation, the general conclusion is that products liability is needed more when the type of activity of the system, or the type of potential damage it may cause, are less covered by insurance than by other systems.¹⁶⁵ But this is true regardless of "thinking algorithms" or of the debate over which tort legal framework ought to apply to "autonomous" systems. What may be said more specifically about "thinking algorithms"? One distinction is that the more foreseeable the system (based on the various parameters discussed above), the more willing insurance companies are likely to be to offer insurance at "reasonable" premiums. This contrasts with algorithms that are less foreseeable, hence are associated with more uncertainty as to the expected amount of damage claims brought against the manufacturers and their expected outcomes, which might render insurance companies reluctant to insure the product at premiums that are not exceedingly high.¹⁶⁶ The division between "more foreseeable" and "less foreseeable" algorithms (again, based on the distinctions made above) in the context of availability of insurance is more relevant when the system is new to the market, and statistics of its performance are not yet available. Once experience is gained as to the system's rates of causing damage, the fact that such a system's choices are not foreseeable in the context of each specific user is not significant. Rather, insurance companies will look at the "big picture" of damage claims brought and won over time. The effect of lack of foreseeability resulting in less available insurance, resulting, in turn, in more compatibility with products liability rationales, is therefore temporary, for a certain initial duration after the product was launched. (For certain systems, especially medical-related ones, the

¹⁶² Polinsky and Shavell, *Supra* Note 82.

¹⁶³ *Ibid.*

¹⁶⁴ *Ibid.*

¹⁶⁵ Damages caused by car crashes, for example, are rarely brought to courts, because such damages are usually covered by automotive insurance. Smith, *Supra* Note 19, 33.

¹⁶⁶ Hubbard, *Supra* Note 86. ("*[T]here may well be no such data available to insurers where a seller seeks liability insurance for an innovative sophisticated robot. As a result, products liability insurers may be very concerned about the potential for high claims. Therefore, insurance may be hard to get, very expensive, or both.*")

duration might be very long considering long-run damages that might appear only years after the user's encounter with the system.)

Secondly, in the absence of full coverage insurance, another major consideration affecting victims' ability to receive compensation is whether an attorney will agree to take their case. Under the contingency fees structure typical of products liability, that likelihood is affected by the attorney's estimation of the probability and magnitude of success, but also of the expected costs.¹⁶⁷ The longer the procedure is expected to last, for instance, due to its technical complexity and the need to hear more expert witnesses, the less likely attorneys will be to take a risk and accept a products liability case.¹⁶⁸

Under the prevalent risk utility test for design defects, a victim must show that efficient safer designs of the system were within reach.¹⁶⁹ This might require the involvement of more than one expert witness, thus rendering the procedure cost-prohibitive.¹⁷⁰ The number of experts needed, hence the complexity, duration and cost of procedures, will probably depend on the system's field of operation, and whether it replaces human professional judgment or not. For example, in a products liability suit against a Roomba vacuum cleaner for injuring a user by catching and pulling her hair,¹⁷¹ a coding expert would be required to show that the manufacturer could have efficiently programmed the system to prevent such accidents. In fields that involve complex professional discretion, such as medicine, engineering or law, it is likely that in addition to a code expert testifying on the availability of technical programming measures, a professional expert in the underlying field (physician, engineer, lawyer, etc.) would also be required to discuss whether such measures were available and accessible in the field of expertise (naturally, defendant too will likely arm itself with all types of relevant experts, in trying to show that no such measures existed). It could be argued, therefore, that systems that replace human professional judgment are more likely to be associated with longer and more expensive procedures (due to the additional experts required but, quite intuitively, also to the complexity of the matter in general) and thus

¹⁶⁷ Smith, *Supra* Note 19, 37-40

¹⁶⁸ Id., Gurney, *Supra* Note 89, 265-266, Hubbard, *Supra* Note 86, 1826-1828.

¹⁶⁹ See discussion in Part II.

¹⁷⁰ Gurney, *Supra* Note 89, 265-266.

¹⁷¹ *Supra* Note 21.

less fulfill the rationale of compensating the victims, as many of the underlying cases will be too expensive to litigate.

Two general characteristics of sophisticated systems that might or might not affect their classification as "thinking algorithm" for the purpose of applying products liability rules are first, the unforeseeable nature of the system (in each of the OODA Loop stages, as discussed above) while statistics of the damages claims associated with said system are not yet available. (The more unforeseeability exists, the more relevant the consideration of compensating the victim becomes because insurance will be less available.) Secondly, if the system replaces professional discretion (in which case the rationale of compensation is less applicable, given that more of these cases will not be picked up by attorneys to begin with). Naturally, both these characteristics may certainly exist at the same time (for instance, systems replacing professional discretion may very likely rely on dynamic sources of information, therefore, as discussed above, be less foreseeable). In such cases they will simply pull in opposite directions with respect to the system's compatibility with the rationale of compensating the victim.

C. Keep it Simple, System

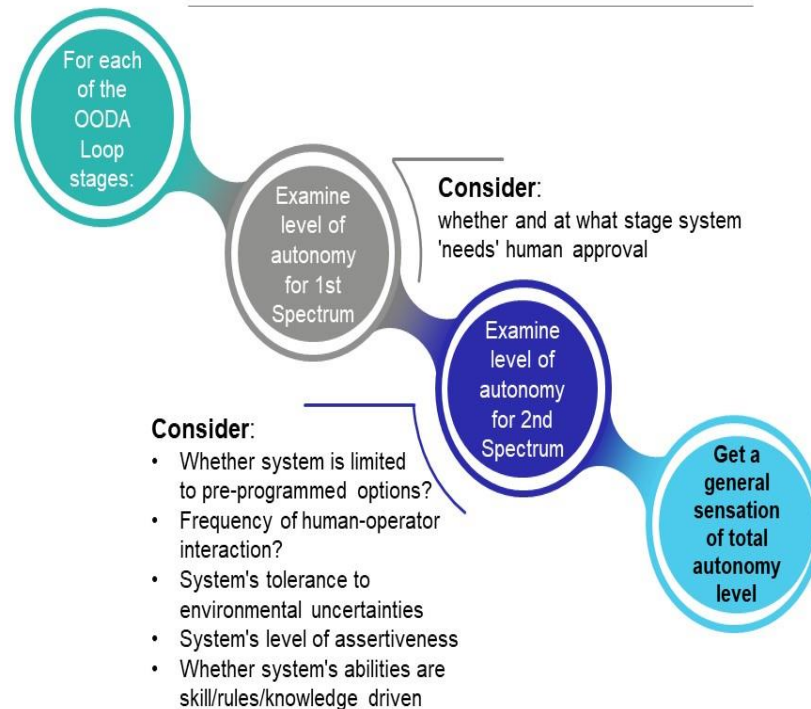
Having analyzed different features affecting systems' compatibility with the rationales behind products liability laws, the article discovered numerous such features that ought to be considered when deciding whether the system is a traditional product or whether a different legal framework should apply to it. Not only is the list of said features relatively long (and open to additional input), its application in practice is not entirely simple, given for instance that certain features might entail more than one effect,¹⁷² or that the effect depends on additional parameters.¹⁷³ As can be seen from the two figures

¹⁷² For instance, when looking at the feature of 'lack of foreseeability and control' as a whole, said feature is not compatible with the rationale of promoting safety (as it might render improvement of safety excessively difficult, costly or inefficient) but might be compatible with the rationale of compensating victims (given that insurance companies would be less willing to insure, thus rendering products liability more needed as an alternative).

¹⁷³ Looking at the example above where lack of foreseeability and control lead to greater need for products liability as an alternative to insurance, said effect is mostly relevant during an initial period of time after the system has been introduced to the market, where statistics of its performance and potential damages are not yet available. Assuming sufficient time has elapsed, the lack of foreseeability and control might be less dominant in the context of insurers' willingness to insure.

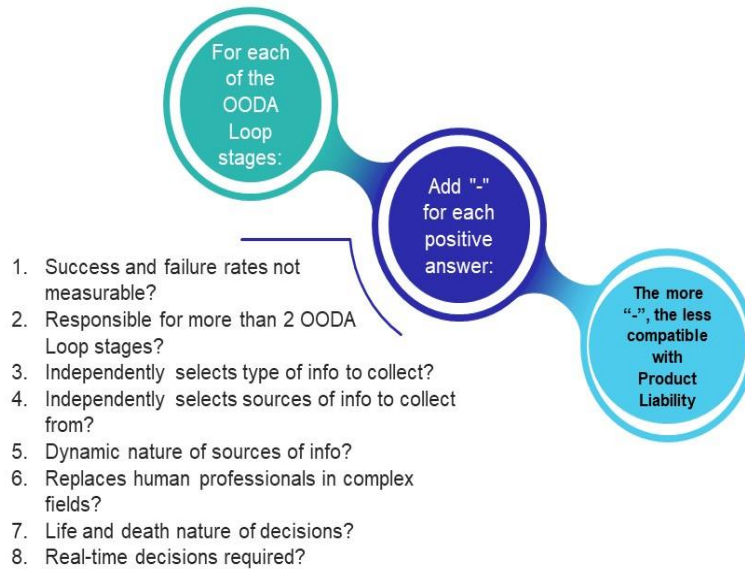
below, however, using said method as a classifier between traditional products and thinking algorithms is nevertheless much more easy and simple than using a system's autonomy level as a distinguishing method between the two.

FIGURE I: AUTONOMY LEVEL AS A CLASSIFIER FOR THINKING ALGORITHMS



Given that lack of foreseeability and control are likely to lead to longer and more expensive legal proceedings, at a certain point in time the dominant effect of this feature on victims' compensation would be limiting accessibility to litigation and thus rendering it again less compatible with products liability.

FIGURE II: PURPOSITIVE INTERPRETATION AS A CLASSIFIER FOR THINKING ALGORITHMS



First, while using autonomy as a classifier requires the user to answer different types of questions, many of which open-ended, and independently decide how to factor in all of her general estimations for these questions, a purposive analysis classifier simply and clearly requires adding a "-" sign whenever a pre-defined feature exists. Secondly, while the outcome received as a result of the former classification process is a general impression of the system's level of autonomy- without practical guidelines as to how to apply it (especially in cases where the result is not definite), the latter classification process yields a clear outcome of the number of "-" signs the system has accumulated. Lastly, said outcome has a practical implication, as it doesn't merely reflect the system's theoretical characteristics, but rather indicates how well the system reconciles with the purposes of the products liability framework.

CONCLUSIONS

Focusing on the legal framework of products liability, the article offered a novel method for determining when damages caused

by 'sophisticated' algorithms ought to be subject to traditional products liability laws, and when an alternative treatment is warranted.

Naturally, preferring a certain alternative legal framework over products liability depends on the characteristics, advantages and disadvantages of such a framework. Proposals for such substitutional options have been discussed for decades,¹⁷⁴ focusing on aspects of determining liability as well as how and by whom damages will be paid. For instance, several analogies have been drawn between sophisticated systems and their creators and the legal relationship between other entities characterized by diminished foreseeability and control over the actions of tortfeasors that are subject to them. Examples were the relationship between parents and minor children who caused damage;¹⁷⁵ between principals and agents,¹⁷⁶ specifically employers and employees;¹⁷⁷ between owners of dangerous animals and their pets;¹⁷⁸ and even between masters and their slaves.¹⁷⁹ Other directions were subjecting the systems themselves to a reasonable analysis in order to determine whether liability by the manufacturer exists,¹⁸⁰ while a different approach focused on no fault insurance schemes to cover the damages caused by such systems.¹⁸¹

A comparison of these different approaches, as well as an assessment of them compared to products liability are beyond the scope of this article. Given the increasing calls to stop treating

¹⁷⁴ See, for example, Sam N. Lehman-Wilzig, *Frankenstein Unbound: Towards a Legal Definition of Artificial Intelligence*, 13 *Futures* 442-457 (1981).

¹⁷⁵ Lehman-Wilzig, *Ibid*; Leon Wein, *The Responsibility of Intelligent Artifacts: Toward an Automated Jurisprudence*, 6 *Harv. J.L. & Tech.* (1992).

¹⁷⁶ Lehman-Wilzig, *Ibid*; Samir Chopra and Laurence F. White, *A LEGAL THEORY FOR AUTONOMOUS ARTIFICIAL AGENTS* (University of Michigan Press, 2011); Peter M. Asaro, *A Body to Kick, but Still no Soul to Damn: Legal Perspectives on Robotics*, in *ROBOT ETHICS: THE ETHICAL AND SOCIAL IMPLICATIONS OF ROBOTICS* 178-180 (Patrick Lin et al. ed, MIT Press).

¹⁷⁷ Lehman-Wilzig, *Ibid*.

¹⁷⁸ Lehman-Wilzig, *Ibid*; Asaro, *Supra* Note 176, 176-178; Richard Kelley et al., *Liability in Robotics: an International Perspective on Robots as Animals*, 24 *Advanced Robotics* 13 (2010); Sophia Duffy and Jamie P. Hopkins, *Sit, Stay, Drive: The Future of Autonomous Car Liability*, 16 *SMU Sci. & Tech. Law. Rev.* 101 (2013).

¹⁷⁹ Lehman-Wilzig, *Ibid*.

¹⁸⁰ Abbott, *Supra* Note 27; Chagal-Feferkorn, *Supra* Note 20.

¹⁸¹ European Parliament Draft Report on Civil Law Rules on Robotics, *Supra* note 22.

"sophisticated" or "autonomous" systems as mere products and subjecting them to a legal framework other than products liability, the focus of this article was to examine why the current classifier between traditional products and systems warranting a new legal framework is insufficient. The article found that said classifier, in the form of the system's level of autonomy, is unduly complex, provides a "fuzzy" and non-practical tool for differentiation, and might yield inconsistent results. The article then proposed a different approach for distinguishing products from non-products for the purpose of applying products liability, focusing on the rationales behind the products liability framework and whether different features of 'sophisticated' systems are compatible with them or hinder them.

Certainly, the parameters discussed in the article warrant additional individual analysis, as they represent general directions of compatibility with products liability rationales, and in specific scenarios might yield different results. The examination of additional types of systems from various sectors employing various technology features too will likely add to the list of relevant parameters that could be of relevance. This article, however, assembled a list of central parameters that are relatively easy to identify and classify as existent or not for each system, and analyzed whether they point to compatibility with the rationales of products liability or not. Decision-makers and scholars encountering a system whose parameters heavily indicate one direction or another could therefore use the proposed analysis to determine whether to apply products liability or to adopt or develop alternative frameworks.