Chapter 4
Contracts

We scanned the skies with rainbow eyes and saw machines of every shape and size … The sun machine is coming down, and we're gonna have a party.

David Bowie, Memory of a Free Festival

Abstract The starting point is the 2005 "World Robotics"-Report of the UN and the Economic Commission for Europe, mainly focusing on "robots of peace" such as environmental robots, surgical robots and edutainment robots. Here, responsibility and legal accountability for the design, construction, supply, and use of robots, are framed as a matter of risk and predictability in contractual obligations. In addition to artificial doctors and cognitive automata such as commercial software-agents, some riskier applications, *e.g.*, ZI agents and unmanned ground vehicles (UGVs), stand for a further set of legal hard cases. The ability of robots to produce, through their own intentional acts, rights and obligations on behalf of humans, suggests distinguishing between robots as tools of human interaction and robots as strict agents in the legal system. However, as a new form of agent in the field of contracts, the increasingly autonomous behaviour of the robot entails the risk that individuals can be financially ruined by the activities of these machines. Whereas the traditional method of accident control via strict liability policies aims to cut back on the scale of the activity, new models of insurance and legal accountability for robots, *e.g.*, the "digital peculium" of robo-traders, illustrate a sounder approach to the contract problem.

U. Pagallo, *The Laws of Robots: Crimes, Contracts, and Torts*, Law, Governance and Technology Series 10, DOI 10.1007/978-94-007-6564-1_4,

At the very beginning, they were cars. As Åke Madesäter stresses in the *Editorial* of the UN World 2005 Robotics report, "the industrial robot was first introduced in the USA in 1961 and the first applications were tested within the car industry in North America" (*op. cit.*, ix). Japanese industry began to implement this technology on a large scale in their car factories in the 1980s, acquiring strategic competitiveness by decreasing costs and increasing the quality of their products. Western car producers learned a hard lesson and followed the Japanese thinking a few years later, installing robots in their factories during the 1990s. Over the past two decades, robots have spread in both the industrial and service fields: as shown by the Report of the Economic Commission for Europe and the International Federation of Robotics (UN World Robotics 2005), we already have "machines of every shape and size," for which the Report provides an analysis on the profitability of robot investments, effects of the business cycle on such investments, the degree of concentration in different countries with prices and wages, the worldwide operational stock of different types of robots, up to the value of the world robot market in the period of 1998–2004. Admittedly, in the extremely dynamic field of robotics, such data becomes quickly out of date. However, this Report allows us to preliminarily understand the panoply of robotics applications with which we are confronted when defining clauses and conditions of contracts.

On one side, we are dealing with a class of industrial robots employed in a number of fields as different as for example, agriculture, hunting, forestry, fishing and mining. These robots are used in the manufacture of food products and beverages, textiles and leather products, wood and coke, rubber, plastic products and basic metals. They are also used when refining petroleum products and nuclear fuel, producing domestic appliances and office equipment, electrical machinery, electronic valves, tubes and other electronic components; as well as semiconductors, radio, television and communication equipment; medical precision, motor vehicles and so on. The properties of these robots can be summarized according to the ISO 8373 definition as "an automatically controlled, reprogrammable, multipurpose manipulator programmable in three or more axes, which may be either fixed in place or mobile for use in industrial automation applications" (UN 2005: 21). The programmed motions or auxiliary functions of these robots can be changed without physical alteration, that is, without the alteration of the mechanical structure or control system except for changes of programming cassettes, ROMs, etc. In connection with the axis or direction used to specify the robot motion in a linear or rotary mode, their mechanical structure suggests a further distinction between Cartesian robots, cylindrical robots, SCARA robots, articulated robots, parallel robots and so forth.

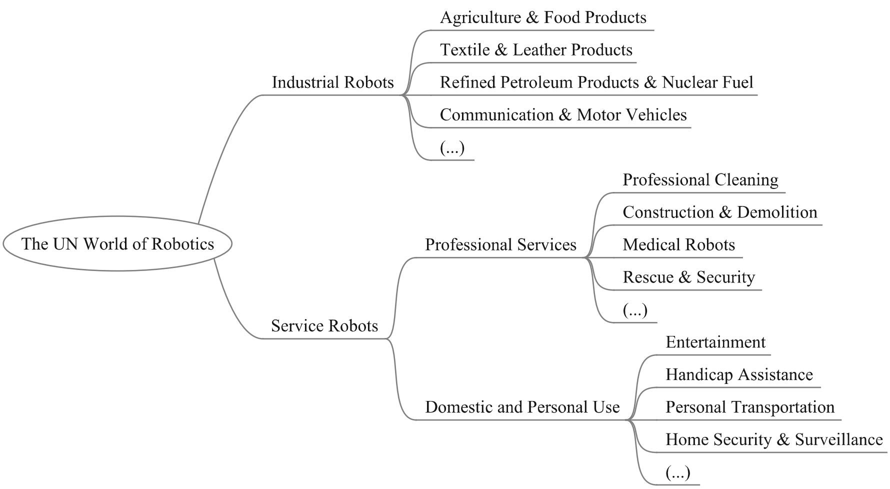

Fig. 4.1 Contractual obligations and robotics complexity

On the other side, we are also addressing a class of service robots that include professional service machines as well as domestic and personal use of robots. In the first subset, we find robots for professional cleaning, inspection systems, construction and demolition, logistics, medical robots, defence, rescue and security applications, underwater systems, mobile platforms in general use, laboratory robots, public relation robots, etc. In the second subset, there is the personal use of robots for domestic tasks such as iRobot's Roomba vacuum cleaning machines; entertainment robots such as toy robots and hobby systems; handicap assistance; personal transportation; home security and surveillance and so on. Whilst further service robots applications should be mentioned, *e.g.*, the new generation of robo-traders examined below in Sect. 4.3, such differentiations are critical in order to discern matters of responsibility and legal accountability for the design, construction, supply and use of robots, through notions of risk, safety, predictability, strict agency, delegation, and so forth, in the civil (as opposed to the criminal) law field. We can begin to chart the complexity of the field according to Fig. 4.1:

What all the robots in Fig. 4.1 have in common is a set of individual rights and obligations that on the basis of voluntary agreements between the parties to a contract concern the design, production and employment of these machines. The aim of this Chapter is to distinguish such voluntary agreements in connection with the level of risk and predictability of robotic behaviour, so as to determine whether basic concepts of contractual law, such as malfunction liability or breach of warranty, are being strained. Clauses and conditions of contracts for the construction and use of robots are differentiated next in Sect. 4.1 in light of a spectrum. At one end, there are a number of reasonable safe and controllable robots; at the other end, we find certain risky applications that can represent ultra-hazardous activities much as traditional aviation was perceived of in the 1930s.

The focus in Sect. 4.2 is on the first end of this spectrum as illustrated by the controlled settings of the operating theatres of the da Vinci surgical robots. Such machines can give rise to engineering problems that scholars routinely address as part of their research, as much as they did with previous technological innovations. On the basis of the probability of events, their consequences and costs, there is a general agreement on how lawyers should define matters of unpredictability and risk as caused by such robots, in order to ascertain individual responsibility for the design, production and use of reasonable safe machines. This class of plain (as opposed to hard) cases refers to notions of evidence, traditional negligence and cases of no-fault responsibility.

The other end of this spectrum, namely certain riskier robotic applications such as the Zero Intelligence ("ZI") agents in the business field, are the focus in Sect. 4.3. The aim is to further distinguish between robots as simple tools of human interaction and robots as proper agents in the civil law field. Although current rules bar the acceptance of the legal agency of robots in certain cases, such legal agency makes sense in that humans delegate relevant cognitive tasks to robots. These machines can send bids, accept offers, request quotes, negotiate deals and even execute contracts, so that the level of autonomy, which is insufficient to hold robots criminally accountable for their behaviour, is arguably sufficient to acknowledge new forms of artificial agency in the law of contracts.

Accordingly, Sect. 4.4 explores new forms of accountability for the behaviour of robots as well as traditional ways of distributing risk through insurance models or authentication systems. The ultimate aim is to avert legislation that makes people think twice before using or even producing robots that provide "services useful to the well-being of humans" (UN World Robotics 2005). The idea that (certain types of) robots may be held directly accountable for their own behaviour has a precedent in the ancient Roman law institution of *peculium*. In Justinian's Digest, the mechanism of *peculium* enabled slaves, deprived of personhood as the ground of individual rights, to act as estate managers, bankers or merchants. Similarly, I suggest that a sort of portfolio for robots could guarantee the rights and obligations entered into by such machines. Drawing a parallel between robots and slaves is attractive, since the aim today is the same as lawyers pursued in Ancient

Rome: individuals should not be ruined by the decisions of their robots and any contractual counterparties of robots should be protected when doing business with them.

After examining surgical robots and cognitive automata in the form of commercial software-agents, or robot-traders, Sect. 4.5 dwells on the case of unmanned vehicles and, more particularly, unmanned ground vehicles such as AI cars and chauffeurs. The reason for this is twofold. On one hand, these kinds of robotic applications allow us to deepen issues of contractual liability and both human and robotic accountability in terms of apportioned responsibility. On the other hand, AI chauffeurs suggest that we will increasingly address (or be pressed by) cases of extra-contractual responsibility, *e.g.*, robots damaging third parties rather than affecting contractual counterparties. In the event a machine fortuitously harms someone in the roundabouts, who shall pay?

4.1 Pacts, Clauses and Risk

Risk can be conceived of in three ways. First, from an evolutionary stance, we can associate the notion of risk with every adaptive attempt to reduce the complexity of the human environment. In their introduction to *Risk Analysis and Society* (2004), Timothy McDaniels and Mitchell J. Small stress that "since the beginning of human development, risks to health and well-being have led to adaptive responses that open paths for change. When Neolithic family groups shared knowledge and resources for combating hunger, thirst, climate, or outside attack, they were trying to manage risks they faced… Risk management has been a fundamental motivation for development of social and governance structures over the last 10,000 years" (*op. cit.*).

A second approach insists on the peculiar features of current modern risk societies and what therefore distinguishes them from traditional (or pre-modern) organizations as well as early modern societies. A classical text such as Ulrich Beck's 1986 *Risikogesellschaft* makes this point clear: "[W]e are eyewitnesses – as subjects and objects – of a break within modernity, which is freeing itself from the contours of the classical industry society and forging a new form – the (industrial) risk society… The argument is that, while in classical industry society the 'logic' of wealth production dominates the 'logic' of risk production, in the risk society this relationship is reversed" (1992 English edition: 9, 14).

A final approach to the notion of risk is methodological: we have to determine the level of risk through quantitative and qualitative evaluations of safety factors, risk assessment and management in terms of probabilities, engineering risks, health risks, information risks and so forth. According to Frank Knight's seminal remarks in *Risk, Uncertainty and Profit* (1921, reissue 2005), we should preliminarily grasp that "risk, as loosely used in everyday speech and in economic discussions, really covers two things which, functionally at least, in their causal relations to the phenomena of economic organization, are categorically different." Those two things are proper risk as "a quantity susceptible of measurement" or "measurable uncertainty," and risk that may be difficult or impossible to quantify, referred to as proper uncertainty. For example, when dealing with the safety factors of structural engineering, *e.g.*, the safety structures of buildings, scholars distinguish between sources of failure amenable to probabilistic assessment, such as poor qualities of materials and higher loads than those foreseen in the project, and uncertain factors such as human error, potentially unknown failure mechanisms, or the imperfect theory of the failure mechanism in question, *i.e.*, proper uncertainty.

Although these three approaches to the notion of risk are intertwined, let us restrict our focus to cases of strict risk and the ways scholars address the challenges of proper uncertainty. A fruitful illustration is offered by the nuclear industry and how, in the 1950s and 1960s, engineers designing nuclear reactors intended to keep the probability of accidents as low as possible, although they did not have any methodology to determine such probabilities. In fact, modern probabilistic risk assessment developed only in the late 1960s and early 1970s, culminating with the 1975 Rasmussen report. In the phrasing of Neelke Doorn and Sven Hansson (2011: 155), "the basic methodology used in this report is still used, with various improvements, both in the nuclear industry and in an increasing number of other industries as a means to calculate and efficiently reduce the probability of accidents." In a nutshell, this probabilistic approach aims to single out the undesirable events to be covered by the analysis, so as to pinpoint the accident sequences that may lead to the occurrence of adverse events as well as the probability of each event in the sequence.

In light of early versions of probabilistic risk assessment, two "improvements" in today's approach should be mentioned. First, experts do not aim at establishing the overall probability of a serious accident but rather, the weaknesses in the safety system, by ranking the accident sequences in connection with the probability of their occurrence. Then, probabilities are not conceived of as "unbiased predictors of occurrence frequencies that can be observed in practice," but "as the best possible expression of the degree of belief in the occurrence of a certain event."[1] This is why, back to the view of

[1] See the definition of the "probabilistic model code" proposed by the Joint Committee on Structural Safety (JCSS 2001: 60).

Doorn and Hansson, experts of probabilistic risk assessment "in the nuclear industry have largely given up the original idea that the outputs of probabilistic analysis of event sequences in nuclear reactors could be interpreted as reasonably accurate probabilities of various types of accidents. Instead, these calculations are used primarily to compare different event sequences and to identify critical elements in these sequences" (Doorn and Hansson 2011: 157).

Such constraints emphasize the critical limits of risk analysis, especially when we are confronted with new and untested technologies and, thus, a lack of data. The empirical basis of probabilistic models necessarily hinges on events that are common enough to let scholars collect data about their occurrence and, yet, probabilities of unusual events may be the most relevant ones in risk analysis. Although further methods have been developed to assign probabilities to rare events, such as extreme value analysis, distribution arbitrariness or boot-strapping methodologies, such approaches may fall short in coping with the unpredictable behaviour of autonomous machines. For example, "boot-strapping techniques still require sufficiently long data records and a careful analysis of the influence of data sampling uncertainties" (Doorn and Hansson 2011: 158). Moreover, certain scholars reckon that measurable risks can hardly be assigned to human reactions vis-à-vis novel or experimental technologies. Rather than hinging on probabilities, the focus should be on qualitative or human-centred approaches, so as to delimit the sphere of uncontrollable uncertainties by singling out new types of human failure (Mosneron-Dupin et al. 1997).

Leaving aside further risk analysis approaches, such as the "partial safety factors" proposed by Isaac Elishakoff (2004), we may wonder how the advancement of robotics technology affects the field. As mentioned in the introduction to this chapter, contractual obligations and rights concerning the design, construction and use of robots are strictly related to the level of risk and predictability of their behaviour. Whereas in *The Laws of Man over Vehicles Unmanned* (2008), Brendan Gogarty and Meredith Hagger claim that "determining fault in complex software and hardware is already difficult" (*op. cit.*, 123), let us consider three different scenarios.

First, we have the da Vinci surgical system that, according to the website of its manufacturer, Intuitive Surgical, "enables surgeons to perform delicate and complex operations" such as prostatectomy procedures, "through a few tiny incisions with increased vision, precision, dexterity and control." Work in the *Mechanical Failure Rate of da Vinci Robot System* shows that only 9 out of 350 procedures (2.6 %) could not be completed due to device malfunctions (Borden et al. 2007). Likewise, in *Device Failures Associated with Patient Injuries During Robot-Assisted Laparoscopic Surgeries* (2008), Andonian et al. affirm that only 4.8 % of the malfunctions that occurred in a

New York urology institute from 2000 to 2007 were related to patient injury. What happens, from a legal viewpoint, in such cases where these artificial doctors do not properly work is examined next in Sect. 4.2.

The second scenario is illustrated by the mishap rate of the unmanned aerial vehicles such as the US Air Force's RQ-1 Predator or the US Army's RQ-2 Pioneer. According to the US Air Force's catalogue, we should distinguish three classes of accidents:

(a) Class A mishaps that include the destruction of $ 1 million in property, loss of a Department of Defence aircraft, or a human casualty resulting in loss of life or permanent disability;
(b) Class B mishaps that involve a $ 200,000–$1 million in property damage, human casualty leading to partial disability or three or more hospitalized personnel; and
(c) Class C mishaps that finally concern a $ 20,000–$ 200,000 in property damage or non-fatal injury leading to a loss of time at work.

By 2005, the level of risk for UAVs was much higher than for traditional aircrafts. When compared to manned aviation, the US Air Force's RQ-1 Predator had 32 times as many accidents per flight-hour, the US Navy's RQ-2 Pioneer more than 300 times and the US Army's RQ-5 Hunter approximately 60 times as many as traditional manned aviation. Accordingly, Peter Singer estimates in *Wired for War* (2009) that notwithstanding technological advancement, training or safer operations under peacetime conditions, UAV security "needs to improve by one to two orders of magnitude to reach the equivalent level of safety of manned aircraft."

Such poor figures certainly characterize the civilian use of UAVs as well. Remarkably, the American National Transportation Safety Board ("NTSB") examined three cases of domestic UAV mishaps between 2006 and 2008. In the wording of Geoffrey Rapp's work on *Unmanned Aerial Exposure* (2009), let us see what occurred in one of these cases:

> In April 2006, a Predator UAV used by the United States Customs and Border Protection Service crashed into the Arizona desert when its operators turned off its engine. When one of the Predator's two ground control stations locked up during flight, its operator switched to the other station but neglected to 'align consoles,' inadvertently cutting off the platform's fuel supply. As the UAV lost power during flight, it began to 'shed electrical equipment to conserve electrical power' [according to the NTSB report].
>
> Although no one on the ground was injured, 'the accident didn't help the industry's reputation' (Stew Magnuson). The UAV glided as close to 100 feet from two homes before striking the ground; homeowners heard the crash and thought a bomb had exploded. The NTSB attributed the crash to inadequate surveillance of the program, pilot error, and inadequate maintenance procedures performed by the manufacturer.

> … Accidents like this have thankfully caused no injuries to date, but widespread use of UAVs in the domestic setting would inevitably produce casualties and property loss as a result of crashes or objects falling from airborne UAVs (G. Rapp, *op. cit.*, 628–629).

The final scenario has to with the point of view of insurance companies and risk management. Such companies are third parties to contracts that either pay out when someone else commits a tort against the insured, or cover losses sustained by the insured against a premium, *i.e.*, the factor through which the sum to be charged for a certain amount of insurance coverage is established. Consider the civilian employment of UAVs and how different uses of such technology are covered by policies such as business or pleasure, commercial or industrial aid. According to Geoffrey Rapp (2009), one commercial UAV imagery company, Moire Inc., "carries $2 million in liability insurance and invites customers to request categorization as 'Additional Insured' under its policy" (*op. cit.*, 647). Moreover, when UAVs are employed for scientific purposes, the premium "has been nearly 85 % of the cost of operation per flight hour" and with respect to hull insurance policies, their cost "has been estimated to reach 2 % of UAV replacement value, plus 0.5 % of ground station replacement value and $30,000 per UAV mission" (*ibid.*).

What these examples of insurance costs suggest is the need to grasp the panoply of robotics applications and their impact on clauses and conditions of contractual obligations in light of a spectrum. At one end, we find a number of reasonable safe and controllable robots that, due to the well-established quantifications of the probability of events, their consequences and costs, do not raise particular challenges to traditional risk assessment analysis or the risk management of insurance companies. At the other end of the spectrum, the progressively unpredictable behaviour of robots raises problems of proper uncertainty, rather than quantifiable risk in the construction and use of these machines. The more we widen the settings and goals of robotic programs, the more we will be dealing with growing amounts of complexity, so that the risks emerging will exponentially increase as a consequence of robotic behaviour. Although we do not have to accept Curtis Karnow's idea that the advancement of robotics will end up in a failure of legal causation as discussed above in Sect. 3.5, it is likely that in the field of contracts, the growing autonomy of robots will affect basic concepts such as foreseeable harm, individual negligence or fault. By considering cases of reasonable safe and controllable robots as seen in the next section, we set the background for the analysis of a new generation of robots that fall within the loopholes of today's legal framework and are further discussed in Sect. 4.3.

4.2 The Artificial Doctor

This section focuses on the case of the da Vinci surgical system as an example of how a significant number of robotic applications do not challenge today's legal framework on matters of liability for the behaviour of such machines. This does not mean, of course, that robotic surgery does not raise certain critical issues. For example, in *Predicting the Long-Term Effects of Human-Robot Interaction* (2011), Edoardo Datteri points out "cases of harmful (occasionally fatal) events brought about by negligent use of medical robots behaving normally." Although da Vinci surgical systems may reduce hospital stays by about one-half and hospital costs by about a one-third, there is the risk of "negligence due to poor training with the robotic system: surgeons [are] not given enough time and resources to learn to use the robot properly, … whereas surgeons with extensive robotic experience declare that it takes a minimum of 200 surgeries to become proficient at the Da Vinci" (*op. cit.*) In *Robotic Surgery Claims on United States Hospital Websites* (2011), Linda Jin et al. argue that the use of such robots appears more as a marketing tool to attract patients than a medical system to improve their care. Through a systematic analysis of 400 randomly selected US hospital websites in June 2010, Jin et al. reckon that "forty-one percent of hospital websites described robotic surgery. Among these, 37 % presented robotic surgery on their homepage, 73 % used manufacturer-provided stock images or text and 33 % linked to a manufacturer website. Statements of clinical superiority were made on 86 % of websites, with 32 % describing improved cancer control and 2 % described a reference group. *No hospital website mentioned risks. Materials provided by hospitals regarding the surgical robot overestimate benefits, largely ignore risks and are strongly influenced by the manufacturer*" (*op. cit.*, italics added). Significantly, the Los Angeles Times published an article on 17 October 2011 by Amber Dance, summing up some of these concerns: "Robotic surgery grows, but so do questions. The Da Vinci system is now in 2,000 hospitals. But there's concern that hands-on surgery still has advantages."

From a legal viewpoint, however, both the design and construction of such robots, as well as their employment in 2,000 hospitals, do not seem particularly challenging. As shown by the case, *Mracek v. Bryn Mawr Hospital*, discussed below in Sect. 4.2.2, the current legal framework concerning liability issues for harm caused by the malfunctioning of electronic devices can properly address harms induced by robotic breakdowns. Yet, such cases of liability do not only have to do with clauses and conditions of contracts established between private persons; namely, in the case of the da

Vinci surgery system, the designer and producer of such robots, Intuitive Surgical, and the user of these machines, such as hospitals and natural (rather than artificial) doctors. In fact, the use of such robots may concern the rights of third parties as well as obligations imposed by the state so as to compensate for any damages done by wrongdoing. Therefore, how clauses and conditions of contracts may involve rights and interests of third parties and, *vice versa*, how the legal protection of third parties may affect contractual rights and obligations are examined in the next Sect. 4.2.1. The focus then is on the claims of a third party, Roland C. Mracek, filing suit against both the producer of the da Vinci surgery system and one of its users, the Bryn Mawr hospital in Philadelphia, due to the malfunctioning of a da Vinci system, as explored in Sect. 4.2.2.

4.2.1 Parties, Counterparties and Third Parties

The employment of robotic applications concerns clauses and conditions established by the parties to a contract as well as the rights and interests of third parties. In addition to insurance companies as third parties covering either losses sustained by the insured or paying off when the insured harms another party, consider what occurred to certain homeowners in the Arizona desert in April 2006. These homeowners heard a Predator UAV gliding as close as 100 ft to their houses before striking the ground and making them think that a bomb had exploded. Luckily no injuries were caused by the UAV.

Two types of obligations must be distinguished concerning designers, producers and users of robots that may damage third parties. Some obligations depend on a voluntary agreement between private persons, others are generally imposed against the will of the agent. This type of extra-contractual responsibility includes cases of intentional wrongdoing, negligence-based liability and strict liability. What common law lawyers sum up with the term of tort, may raise forms of apportioned responsibility between the parties to a contract as discussed above in Sect. 2.2.

Let us now view how this complex set of notions works in practice by taking into account a prostatectomy operation by the da Vinci robot. For example, in *Mracek v. Bryn Mawr Hospital*, we have to distinguish four levels of analysis:

(a) The parties to the contract, that is, Intuitive Surgical and the Bryn Mawr Hospital, that determine the conditions for the use (and maintenance) of a da Vinci surgery system;

(b) The insurance company as a third party to that contract on a voluntary basis (although we will examine cases of compulsory insurance in the next section);
(c) Another third party who voluntarily underwent surgery with the da Vinci system, namely, the patient Roland Mracek and his contract with the Bryn Mawr hospital; and
(d) A tort liability suit filed by the patient as Mracek claims to have suffered unwarranted damages caused by both the parties to the contract (*sub a*), that is, Intuitive Surgical and the Bryn Mawr Hospital.

Contractual parties, when establishing the clauses and conditions of their agreement (*sub a*), will thus have to pay attention to the obligations imposed by the state in order to compensate for unjust damages (*sub d*). Consider contracts of software developers that often establish clauses of strong liability limitations and even exemptions for damages caused by their products. *Vice versa*, reflect on the case of US federal contractors that pursuant to 28 U.S.C. § 2671, know that clauses of immunity that protect their contractual counterparties do not extend to them as seen above in Sect. 3.5. In *Mracek v. Bryn Mawr Hospital*, it is noteworthy that one of the defendants, the Bryn Mawr hospital, was dismissed from the suit by court order. Only Intuitive Surgical, the designer and producer of the robot, had to defend itself by showing that the da Vinci robot did not cause any unjust damage. In order to understand how claims of third parties (*sub d*) may affect conditions and clauses of contracts (*sub a*), we shall focus on the different ways the apportioned liability between the parties to a contract depends on three types of extra-contractual responsibility.

First, liability can be ascribed to the tortfeasor for wrongful conduct because that person intended to do harm. Contemplate the case of a doctor who voluntarily causes harm to a patient through the use of the da Vinci robot system. Whereas, in criminal law, the hypothetical of an intentional tort brings us back to the second step of the phenomenology of *Picciotto Roboto*, see above in Sect. 3.4.2, in the civil (as opposed to the criminal) law field, such a wrongful intention severs the link between claims of extra-contractual liability (*sub d*) and previous contractual obligations (*sub a*). It is clear that the producer of the robot is not to be held liable for the conduct of the user of the machine.

Second, there is the opposite case of strict liability, or liability without fault, invoked when the conduct of the tortfeasor is not blameworthy. Regardless of the absence of any illicit or culpable behaviour, individuals are held liable for damages caused by their own dangerous activities or the behaviour of other agents in the legal system. In the case of strict product liability, it follows that claims of extra-contractual responsibility (*sub d*) can overrule contractual agreements for the design, construction and supply of such a product (*sub a*). At times, the producer, rather than the user, of the robot will have to show that there is no evidence that the machine did not properly work.

Finally, liability can be based on negligence or lack of due care, *e.g.*, when a reasonable person fails to guard against foreseeable harm. As mentioned above in Sect. 3.5, strict liability rules do not prevent additional individual liability for careless conduct. Furthermore, a negligence claim may stand even in the absence of a defect under strict liability norms. The link between extra-contractual liability (*sub d*) and contractual obligations (*sub a*) hinges, therefore, on the circumstances of the case, so as to determine whether the user or the producer was negligent.

In light of this general framework, let us deepen how robotic applications affect clauses of civil (as opposed to criminal) responsibility. In this context, we can set aside cases of intentional torts as well as crimes of intent: as shown by the second step of the phenomenology of *Picciotto Roboto* in Sect. 3.4.2, these hypotheticals end up in the class of plain cases. The focus rather should be on strict liability rules and cases of negligence in the civil law field and how the burden of proof is allocated in such cases. Regardless of the differences between common and civil law systems, discussed more thoroughly below in Sects. 5.2 and 3, this complex set of notions and procedures can be illustrated with *Mracek v. Bryn Mawr Hospital*. In this case, the patient/plaintiff alleged that the da Vinci robot caused damage arising out of strict product and malfunction liability, negligence and breach of warranty. The reasons why plaintiff finally lost his case introduce a new class of plain cases in the laws of robots. The general agreement depends on the fact that there are a number of reasonably safe and controllable robots out there.

4.2.2 Producers, Users and Patients

Something went wrong with the surgical removal of a part of Roland Mracek's prostate at the Bryn Mawr Hospital in Philadelphia on 9 June 2005. According to the plaintiff, liability for erectile dysfunction and groin pain following from the medical procedure should be imposed on both the producer (Intuitive Surgical) and the user (Bryn Mawr Hospital) of the da Vinci surgery system. Such a machine would have caused damages, first of all, due to its own malfunctioning, so that the producer of the robot should be held strictly liable. In the phrasing of § 402A of the Restatement (Second) of Torts in the US, strict liability is imposed "not only for injuries caused by the defective manufacture of products, but also for injuries caused by defects in their design." In such cases, the burden of proof falls on the plaintiff who has to prove that the product was defective; that such defect existed while the product was under the manufacturer's control; and, moreover, the defect was the proximate cause of the injuries suffered by the plaintiff. Both the standards and burdens of proof required by § 402A of the Restatement (Second) of Torts apply to liability claims for breach of warranty as well.

Plaintiff's second claim has to do with provisions of strict malfunction liability. Responsibility can be imposed although the plaintiff is not able to produce direct evidence on the defective condition of the product or the precise nature of the product's defect. Rather, the plaintiff is to demonstrate that defect through circumstantial evidence of the occurrence of a malfunction, or through evidence eliminating both abnormal use of the product and reasonably secondary causes for the accident.

Finally, responsibility for civil (as opposed to criminal) negligence hinges on the duty to conform to a certain standard of conduct. Here, the plaintiff has to prove that defendants breached that duty, thereby provoking an injury and an actual loss or damage to the plaintiff.

Interestingly, Mracek did not submit any expert report to support or corroborate his claims. In the wording of the District Court, the plaintiff's argument was that the asserted defect of the robot was "obvious enough to be ascertainable by the average juror without speculation." More particularly,

> Mracek contends that an expert report is not necessary because the surgeon who performed his operation, Dr. McGinnis, will testify at trial concerning not only his pre- and postoperative medical condition, but also the malfunction of the da Vinci robot. Mracek maintains that the defect of the surgical robot is obvious because all of its component parts shut down after repeatedly flashing "error" messages, and then was not able to be restarted once the surgery commenced. Mracek argues that it is not necessary for him to produce an expert report for a finding of an obvious defect, as such a defect is not beyond the purview of a layperson when presented with this factual record (District Court of Philadelphia, Judge R. Kelly, *case 08-296* from March 11, 2009, *cit.*, 6).

Although "absence of expert testimony is not fatal to a products liability case," this principle does not typically apply to such complex machines as the da Vinci robot. All in all, this is why Mracek lost the case. According to the court, the plaintiff failed to support his case without an expert report, because he could not establish either a defect of the robot or a causal link between the problems with the robot and the plaintiff's damages under strict liability rules. Likewise, under the malfunction theory of strict products liability, the plaintiff did not offer any evidence so as to eliminate reasonable secondary causes, nor did he produce any genuine issue of material fact regarding elements of negligence that could be given to a jury. Therefore, the court granted the defendant's motion for summary judgment against Mracek in 2009. Under US Federal Rule of Civil

Procedure 56(c), summary judgment is to be granted "if there is no genuine issue as to any material fact and the moving party is entitled to judgement as a matter of law."

The Court of Appeals confirmed the District Court's judgment in 2010, with Justices Scirica, Barry and Smith rejecting Mracek's argument that the District Court improperly granted summary judgment on his strict malfunction liability claim.[2] The court reasoned that the trial court's decision "was proper because he [Mracek] failed to demonstrate a genuine dispute of material fact. Most importantly, there is no record evidence that would permit a jury to infer Mracek's erectile dysfunction and groin pain were caused by the robot's alleged malfunction" (*op. cit.*, 5). As the plaintiff cannot depend upon simple conjecture or guesswork and has to introduce "evidence from which a rational finder of fact could find in his favour," the Court of Appeals confirmed the trial court's grant of summary judgment. Four months later, Mracek filed a petition for a writ of certiorari before the Supreme Court, which was distributed for conference in September, and a few days later, on 4 October 2010, denied.

After the set of plain cases on crimes of intent as examined in Sect. 3.4.2, the *Mracek vs. Bryn Mawr Hospital* case illustrates a further class of "general agreement in judgments as to the applicability of the classifying terms" (Hart 1994: 123). On one hand, Mracek's case seems plain because of the lack of evidence. On the other hand, as a genuine dispute of material fact, in the phrasing of the Court of Appeals, we can imagine an alternative outcome of the case, *i.e.*, the plaintiff could prove the causal link between the behaviour of the robot and his erectile dysfunction. Yet, traditional notions of the law, such as proximate or reasonable secondary causes, negligence or breach of warranty, would still be at work. The reason why the behaviour of the da Vinci system does not affect how lawyers grasp individual liability in these cases, hinges on the controlled settings of the operational theatres that delimit the conduct of the machine: its mechanisms and properties do not look more intricate than the complex analysis of scientific experts in other fields of the law as raised above in Sect. 3.5. After crimes and torts, both of which depend on the "wrongful" conduct of humans, this class of plain cases referring to hypotheticals of strict malfunction rather than strict product liability,

[2]Mracek's appeal did not concern his previous claims on strict product liability, negligence and breach of warranty. In *Unmanned Vehicles and US Product Liability Law* (2012), Stephen S. Wu addresses further cases where "defendants were entitled to summary judgement because the plaintiffs failed to introduce evidence in opposition to summary judgement showing that the system was defective." Among such cases, see *Jones v. W + M Automation*, 818 N.Y.S. 2d 396 (App. Div. 2006), appeal denied, 862 N.E. 2d 790 (N.Y. 2007); and *Payne v. AAB Flexible Automation*, 96–2248, 1997 WL 311586 (8th Cir. Jun. 9, 1997).

represents the first end of the spectrum of robotic applications, namely, machines that are reasonably safe and controllable.

However, it is not so difficult to conceive of more complex cases. Let us dwell on the Bryn Mawr Hospital and imagine the more than realistic scenario of an artificial agent working at that hospital, scheduling the appointments of patients. The agent checks priorities for surgeries performed by the da Vinci surgery system and alerts maintenance staff and so forth. This robot suggests we are dealing with a proper agent, rather than a simple tool of human interaction. There are already, after all, a number of such agents that terminate or renew Medicaid programs, food stamps and other welfare schemes, by enrolling "applicants directly into benefits programs without review or critique by human operators" (Chopra and White 2011: 195). Furthermore, by widening the set of parameters and conditions regulating the behaviour of the robot, *e.g.*, machines operating in open environments, it is likely that the level of risk and proper uncertainty arising from the use of such machines will severely impact basic tenets of the law and, more particularly, the field of contracts. In *Agent Technology: Computing as Interaction* (2005), Michael Luck et al. draw attention to a number of possible candidates for a new generation of legal hard cases, such as "simulation and training applications in defence domains; network managements in utilities networks; user interface and local interaction management in telecommunication networks; schedule planning and optimisation in logistics and supply-chain management; control system management in industrial plants," up to simulation modelling "to guide decision makers in public policy domains" (*op. cit.*, 50).

Here, the legal challenges of robotics in the field of contracts can be illustrated with a class of machines that may negotiate deals, accept bids, send offers and establish rights and duties of their own. Contrary to the controlled settings of the da Vinci system, the class of trading artificial agents may affect basic notions and ways of legal reasoning in three different ways. First, such machines can successfully be used to carry out complex business transactions and, yet, their behaviour, at times, suggests troubling parallels with the greediness of human speculators. Second, these robots are traditionally presented as instruments of human interaction and, still, an increasing number of scholars reckon that such robots should be conceived as new actors in today's legal systems. Finally, strict liability rules currently apply to robots and, nevertheless, such artificial agents suggest new forms of accountability and responsibility for the behaviour of others in both contracts and tort law. Therefore, let us proceed with the analysis at the opposite end of the spectrum represented by the reasonable safe and controllable robot examined in this section. Matters of risk and, moreover, of proper uncertainty as at the other end of the spectrum, are at stake with a new generation of robo-traders.

4.3 Robo-Traders

Work in artificial trading agents has been cutting edge in the past few years. Along with contributions to the trading agent competition ("TAC")-context, such as Seong Jae Lee et al. in *RoxyBot-06: An (SAA)2 TAC Travel Agent* (2007), we can mention the works of Jeffrey Mackie-Mason and Michael Wellman in *Automated Markets and Trading Agents* (2006), Michael Wellman, Amy Greenwald and Peter Stone in *Autonomous Bidding Agents* (2007), Giovanni Sartor in *Cognitive Automata and the Law* (2009), Samir Chopra and Laurence White in *A Theory for Autonomous Artificial Agents* (2011). Whereas, most of the time, these works focus on software agents, rather than robots interacting in the real world, such machines raise some common issues. On one hand, their behaviour and decisions can be unpredictable and risky, as shown by robotic experiments in double auction markets throughout the past decades. Here, the traditional legal viewpoint considers robots simply as tools or means of human interaction, which means that strict liability rules apply to humans as principals of the machine. On the other hand, there is a number of reasons why some of these robots should be deemed as proper agents rather than tools of human interaction: such machines can be extremely efficient in establishing rights and obligations between humans that delegate to them complex cognitive tasks. As a result, today's strict liability rules raise the threat that people think twice before employing robots that may provide "services useful to the well-being of humans" (UN World Robotics 2005). Richard Posner summarizes this popular stance when claiming that the best method of accident control is to scale back the activity (Posner 1973: 180).

This section sheds light on the legal challenges of robotics through a case study in the field of artificial trading agents. Next, attention is paid to the robotic experiments in double auction markets in Sect. 4.3.1 in order to illustrate the pros and cons of such technological applications. The first laboratory double auction in markets, where buyers and sellers submit bids and offers in any order, was reported by Vernon Smith's classic paper *An Experimental Study of Competitive Market Behaviour* (1962). Some thirty years later, robot tournaments were conducted at the Santa Fe Institute and, in the early 2000s, an Automated Trading project in robots, trading in auction markets, was sponsored by the University of Pennsylvania and Lehman Brothers. This case study is deepened in Sect. 4.3.2: the focus is on the traditional legal viewpoint that holds individuals responsible for the use of such robot traders according to the rules that apply to users as principals of these machines. By showing how today's strict liability rules fall short in coping with certain legal challenges of robot traders in Sect. 4.3.3, the aim of Sect. 4.4 is to provide a more fruitful guide to a new generation of hard cases in the legal domain.

4.3.1 Artificial Greediness

The baseline for all robot archetypes in double auction markets is given by the Zero Intelligence ("ZI") agents. These robots are rudimentary in that they are oblivious to their environment and do not control the timing of their actions: ZI agents even lack the capability of taking action so as to compensate for their inability to respond to the environment. As Ross Miller argues in his telling article *Don't Let Your Robots Grow Up to Be Traders* (2008), a ZI agent is a robot programmed to simply "generate bids and offers selected randomly from a uniform distribution subject only to the constraint it cannot 'deliberately' lose money." However, if ZI agents are certainly rudimentary, they also achieve sophisticated goals as outperforming untrained human traders in double auction experiments. Moreover, the performance of ZI agents in shopping around or planning ahead can be improved, so that according to Miller, "the design of a special-purpose agent that can trade in the simple asset markets… as well as, if not better than, humans seems clearly within grasp" (*op. cit.*).

Interestingly, since the robot tournaments at the Santa Fe Institute in 1990, scholars have programmed ZI agents in order to replicate human double-oral auctions, showing that markets populated only by such robots have the tendency of human markets to generate average prices and quantities of what economists traditionally present as a "competitive equilibrium." As Shyam Sunder affirms in *Markets as Artefacts* (2004), computer simulations have demonstrated "that allocative efficiency – a key characteristic of market outcomes – is largely independent of variations in individual behaviour under classical conditions." This ability of ZI agents to achieve a high level of allocative efficiency when determining average prices and quantities of goods exchanged in a market can be grasped with Friedrich Hayek's idea that in certain fields of social interaction, such as pacts and contractual obligations, "intelligence" emerges from the rules of the game rather than individual choices. Yet, a lot of problems arise when addressing the subtleties of markets containing intelligent agents such as humans. Work on robot trading in auction markets as the Automated Trading project, sponsored by the University of Pennsylvania and Lehman Brothers, showed relevant failures as to programming robot traders capable of effectively speculating against (smart) humans. It is noteworthy that this project was finally suspended in 2005, that is, 3 years before Lehman Brothers' own collapse…

In addition, the complexity of tackling multiple actions occurring synchronically in time far exceeds the capabilities of ZI agents. This circumstance reduces the allocative efficiency of the market and leads to a rudimentary bubble and crash scenario, where traders act without regard of the effects of future supply. As in real life bubbles, agents are overwhelmed by the complexity of the environment, thereby appearing extremely inexperienced. This analogy has suggested that experiments with the random-bidding strategy employed by such robots can clarify how real life bubbles form. As stressed by Miller (2008), "the bubble in Internet and other technology stocks that formed at the end of the 1990s may have been partially rooted in market participants' inability to properly anticipate the future supply of stock in Internet companies." Similarly, others argue "that some of the financial troubles of late 2009 may have been caused by the involvement of such agents operating without human supervision and at speeds not amenable to human understanding or intervention" (Chopra and White 2011: 7).

The parallel between the greediness of human speculators and the eagerness of ZI robots to trade, however, does not mean that such artificial agents should not be preferred to humans in certain market operations, *e.g.*, when speed is valued over intelligence. Moreover, there are a number of robotics applications and, generally speaking, of autonomous artificial agents that do not raise such a level of risk when, say, individuals bid, buy or book. Suffice it to mention today's routine interaction with eBay bidding agents, iTunes store agents, Amazon's website bots, or the common airline booking system that through "yield management techniques," determine prices according to how crowded the flight is and so forth. By opening up new ways of "making business as usual," *e.g.*, granting authority to the artificial agent so as to let it act on an individual's behalf when dealing with third parties, we should pay attention to how the law aims to govern such business. For example, we may agree with the American Law Institute and Commissioners of the Uniform State Laws that contracts made by electronic agents should be considered valid, although no action or knowledge of any human being may be involved. Still, this approach leaves open the question of whether humans are bound by every decision of a robot and which human party would be bound by such decision: the designer/implementer of the robot, its user, the operator or the principal?

4.3.2 The Robot and the Principal

Rights and obligations established by robots can be interpreted through the traditional legal viewpoint as examined already with the artificial doctor. Strict liability rules should in fact govern the behaviour of robots, binding those humans on whose behalf they act, regardless of whether such conduct was planned or envisaged. In the US, for example, the E-SIGN statute and the 1999 attempt to amend the Uniform Commercial Code with a Uniform

Computer Information Transactions Act ("UCITA") illustrate this approach. On the one hand, 15 U.S.C. § 7001(h) provides that a contract "may not be denied legal effect, validity or enforceability solely because its formation, creation or delivery involved the action of one or more electronic agent so long as the action of any such electronic agent is legally attributable to the person to be bound." On this basis, in *Spiders and Crawlers and Bots* (2002), Jeffrey Rosenberg claims that "a robot that enters into a clickwrap agreement, either by clicking on an 'I accept' button, or disregarding the express protocol set forth in a robot exclusion header, binds the person who designed and implemented the robot."

On the other hand, section 107 (d) of UCITA establishes that "a person that uses an electronic agent that it has selected for making an authentication, performance or agreement, including manifestation of assent, is bound by the operations of the electronic agent, even if no individual was aware or reviewed the agent's operations or the results of the operation." Likewise, the Unicitral document enclosed in the proposal of the UN Convention Electronic Communications in International Contracts Documents states that "general principles of agency law (for example, principles involving limitation of liability as a result of the faulty behaviour of the agent) could not be used in connection with the operation of such systems. The Working Group reiterated its earlier understanding that, as a general principle, the person (whether a natural person or a legal entity) on whose behalf a computer was programmed should ultimately be responsible for any message generated by the machine… As a general rule, the employer of a tool is responsible for the results obtained by the use of that tool since the tool has no independent volition of its own."

Summing up the outcomes of the robots-as-tools approach, we consequently have:

(a) Robot *R* acting on behalf of the principal *P*, so as to negotiate and make a contract with the counterparty *C*;
(b) Rights and obligations established by *R* directly bind *P*, since all the acts of *R* are considered as acts of *P*;
(c) *P* cannot evade liability by claiming either she did not intend to conclude such a contract or *R* made a decisive mistake;
(d) In case of the erratic behaviour of *R*, *P* may claim damages against the designer and producer of *R*. However, according to the mechanism of the burden of proof, *P* will have to demonstrate that *R* was defective and that such defect existed while *R* was under the manufacturer's control; and, moreover, the defect was the proximate cause of the injuries suffered by *P*.

Although the traditional outlook may fit under certain circumstances, the robots-as-tools approach is flawed for three reasons. First, it is likely that most of the time, humans will delegate to autonomous and even smart robots complex cognitive tasks, such as acquiring knowledge for decision-making. Consequently, it is difficult to accept the traditional idea that robots are mere tools of human interaction and, moreover, that rights and obligations established by robots would be directly conferred upon humans (*sub b*), because the principal wanted the specific content, or agreement, of the contract made by the artificial agent. Rather, rights and obligations are conferred onto humans because they delegate to the robot the authority to act on their behalf.

Second, from the fact that *P* delegates to *R* (*sub a*), it does not follow that the legal effects of the behaviour of *R* should necessarily fall upon *P* (*sub b*). Admittedly, the robot's counterparty *C* should be allowed to expect, in good faith, that the machine really means what it declares, *e.g.*, a contractual offer, when negotiating with robot *R*, so that *P* cannot evade liability by claiming she did not intend to conclude such a contract (*sub a*). However, humans should not be able to avoid the usual consequence of robots making a decisive mistake, *i.e.*, the annulment of a contract, when *C* had to have been aware of a mistake that due to the erratic behaviour of the robot, clearly concerned key elements of the agreement, such as the market price of the item or the substance of the subject-matter of that contract. Here, it seems reasonable to expect that the humans involved in such transactions should be bound by the interpretation of the behaviour of the robot that usually applies to the circumstances of the case according to existing conventions of business and civil law.

Third, the robots-as-tools approach appears unsatisfactory when responsibility (and risk) must be distributed between, say, operators and users as principals of the robot. Whereas the traditional approach ends up in a Hegelian night where all kinds of responsibility look grey, operators and users of robots should be held accountable in accordance with the different errors of the machine and the circumstances of the case. In fact, the erratic behaviour of the robot can concern not only software and hardware malfunctioning, or errors of specification as mentioned above, *e.g.*, errors concerning the substance matter of a contract. In the phrasing of Chopra and White (2011: 46), we should take into account "induction errors, where a discretionary agent incorrectly induces from contracts where the principal has no objections to a contract the principal does object to." Aside from a further hypothetical of liability involving the manufacturers of the artificial agent, we should also distinguish cases where operators and users of the robot coincide and cases where operators allow users to use the machine, so as to deal with third parties. Nine possible cases follow as a result: the legal variables of this section are illustrated with Table 4.1. "Yes" and "no" refer

Table 4.1 What the approach to robots-as-tools lacks

Erratic robot	Specification	Induction	Malfunction
Human operator	Yes	Yes	Sometimes no
Human user	Yes	No	Sometimes no
Third parties	No	No	Sometimes yes

to whether or not human operators, users or third parties should be held accountable for the erratic conduct of the machine:

In *A Legal Theory for Autonomous Artificial Agents* (2011), Chopra and White examine this complex scenario by further considering the theories of the unilateral offer, of the objective intention, and so forth (*op. cit.*, 45–50). Here, it suffices to pay attention to the three rows of Table 4.1. The first set of cases concern legal responsibility of the human operator for the erratic behaviour of the robot due to specification errors, induction mistakes or the malfunction of the machine. Compared with the strict liability approach, according to which operators might be liable under all circumstances, it is arguable that such an operator should not be accountable for malfunctions of the machine that are obvious to users and third parties. In the wording of *A Legal Theory for AAAs*:

> An example of the first kind of transaction occurs when the principal is the operator of a shopping website (such as Amazon.com), the agent is the website interface and backend, and the third party is a user shopping on the website. The contract is formed between the principal and the third party…
>
> When the principal is the agent's operator, specification and induction errors will be less obvious to third parties than to principal/operators, and therefore the principal/operator will normally be the least-cost avoider of the loss. Where, for example, because of specification or induction error, a book is advertised very cheaply, the third party may simply understand the price to be a "loss leader" rather than the result of an error… In the case of malfunction it may be obvious to the third party, because of other indications, that a particular price is the result of error… Therefore, often, the least-cost avoider of malfunction errors will be the third party.
>
> With the agent understood as a mere tool, the principal would be liable for all three types of error in all cases. This approach would not be efficient where the third party is the least-cost avoider of the risk, as in many cases of malfunction error (Chopra and White, *op. cit.*, 46–47).

Vice versa, we can imagine cases where the principal is the user, rather than the operator, of the artificial agent. After all, this is what occurs on eBay, where individuals use the auction website's proxy bidding system so as to enter a contract with a third party:

> In this case, as in the operator as principal case, the risk of specification errors should normally fall on the principal, that is, the user of the agent. However, the risk of induction errors should normally fall on the operator of the agent (who has control over the agent's design and operation). The risk of malfunction errors will

often most fairly fall on the third party, for the reasons given in discussing the operator as principal case.

Under the "agent as mere tool" solution, the user/principal would be primarily liable for all three types of error, incorrectly allocating the risk of induction and malfunctions error in particular (Chopra and White, *op. cit.*, 48–49).

The final row of Table 4.1 concerns responsibility of third parties for the threefold erratic behaviour of robots. As stated in this section, the traditional legal stance falls short in coping with the accountability of those who have to be aware of, say, a mistake of the robot due to its erratic behaviour. Aside from the allocative efficiency of such no-fault responsibility rules, there is the risk that strict liability policies can dissuade humans from employing robots at all. Is there a feasible way out of the *cul-de-sac* that characterizes the robots-as-tools approach?

4.3.3 A New Agent in Town

It makes a lot of sense to conceive (certain types of) robots as proper agents in the field of contracts, that is, granting them the authority to act on an individual's behalf when dealing with third parties. Such a perspective prevents certain key flaws of the robots-as-tools approach, since the legal agency of the robots makes it clear that humans do delegate crucial cognitive tasks to these machines. We can establish individual responsibility for the erratic behaviour of robots properly, taking into account the "intentions" of such machines and moreover, by referring them to existing conventions of business and civil law. As stressed in Sect. 4.3.1, we should take the idea seriously that robots have intentions relevant in the civil (as opposed to the criminal) law, for intelligence emerges from the rules of the contractual game, rather than individual choices of the robotic agent. In the phrasing of Giovanni Sartor:

> [T]his leads to assimilate the situation of the user of [a robot] to the situation of a person handing over the conclusion of a contract to a human agent… What the two situations have in common, which distinguishes them from the situation where one uses a (mechanical or human) means of transmission, is cognitive delegation, *i.e.*, the decision to entrust the formation of the content of a contract and the decision whether to conclude it or not… to someone (or something) else's cognition (Sartor 2009: 280–281).

Admittedly, the current rules of legal systems bar the acceptance of the robots-as-agents approach in certain cases. Furthermore, there are key differences as to how common and civil law systems may aim to govern such technological applications. For example, in France or Italy, the legal personality of the agent is a necessary (yet not sufficient) requirement for acknowledging that machines can be proper agents in the civil (as opposed to the criminal) law field. *Vice versa*, in Anglo-American law, there is no objection "to the possibility of a nonperson artificial agent, on the grounds of a lack of capacity to contract in its own right on the part of the agent" (Chopra and White 2011: 56). Likewise, in the US, the principal is not bound by a contract that is outside the agent's actual or apparent authority, although a "minimum of physical and mental ability" or "volition" of the agent is required. In most civil (as opposed to common) law systems, the agent must be of sound mind, so that the risk of malfunction errors would fall on the third parties in all cases. However, despite this general disagreement, we should not overlook a crucial point: robots should be conceived as new proper agents in the civil law field because this legal option allows us to strike a fair balance between the individual's claim to not be ruined by the decisions of their robots and the claim of a robot's counterparty to be protected when doing business with them. Some brief remarks on the history of the law help us in the next Section: Roman lawyers addressed both legal agency of non-humans and guarantees for the counterparties interacting with them more than 2,000 years ago. A historical reference on the rules that governed the actions of slaves sheds light on how we could deal with today's robots following the pragmatic spirit of Roman law. The analysis of the ethical issues raised by this parallel, is postponed until Sect. 6.1.

4.4 Modern Robots, Ancient Slaves

The parallel between today's robots and slaves in ancient Rome seems appropriate, because slaves were considered as things that nevertheless played a crucial role in trade and commerce. In *The Human Use of Human Beings* (1950), the father of cybernetics, Norbert Wiener, suggested that "the automatic machine, whatever we may think of any feelings it may have or may not have, is the precise equivalent of slave labor." This similarity has been stressed time and again over the past years. In *The Responsibility of Intelligent Artifacts* (1992), Leon Wein reckons that automation is "bringing the conception of slavery back on the scene… As employees who replaced slaves are themselves replaced by mechanical 'slaves,' the 'employer' of a computerized system may once again be held liable for injury caused by his property in the same way that she would have if the damage had been caused by a human slave" (*op. cit.*, 111).

From a legal viewpoint, however, we should not miss the forms of agency that ancient Roman law admitted for such "things." Although most slaves certainly had no rights to claim against their own masters, some of them enjoyed a significant autonomy. The elite of the slaves, as in the case of the emperor's slaves, were estate managers, bankers and merchants, holding important jobs as public servants, or entering into binding contracts, managing and making use of property for their masters' family business. Consider the case of the *institor* (*Dig*. XIV, 3, 11, 3; XV, 1, 47). Such slaves managed different classes of convenience stores, *taverna*, such as bakeries and barbershops; wineries, hot drinks, or ready-prepared meat; and even, so to speak, booksellers' minimarts. When Emperor Nero was convinced to participate in the Olympic games of 67 A.D. in order to improve relations with Greece, it was not a joke that he entrusted his freedman Helios with the right to convict or seize anyone in Rome.

The parallel between robots and slaves is hence attractive, because the rules of ancient Roman law on slavery show a way to address certain of the inconsistencies of the robots-as-tools approach mentioned in the previous section. While Roman lawyers invented forms of agency and autonomy for mere things without legal personality, their aim was to strike a balance between the interest of the masters not to be negatively affected by the business of their slaves and the claim of the slaves' counterparties to be able to safely interact or do business with them. Today's idea that (certain types of) robots may be held directly accountable for their own behaviour has thus a precedent in the ancient Roman legal mechanism of *peculium*. In order to avert any legislation preventing the use of robots due to excessive burdens on the owners (rather than producers and designers) of these machines, the idea is that, at times, only "robots shall pay" could be the right answer.

4.4.1 The Digital Peculium

There is a key difference between criminal and civil lawyers dealing with new types of responsibility for the behaviour of robots. The focus of criminal lawyers is most of the time on harm or damages caused by such machines: something had to go wrong, in other words, so as to determine whether we are dealing with crimes of intent, negligence, or further legal observables examined with the phenomenology of *Picciotto Roboto* in the previous Chapter. *Vice versa*, it is not necessary that something has to go wrong in civil law: on the contrary, since the late nineteenth century, the legal imagination has been fired by how machines can be extremely fruitful in making contracts, or establishing rights and obligations between humans, in a win-win scenario. Although today's debate on cognitive automata in the form of software agents can be traced back to the seminal remarks of German scholars on automation and the law in the late 1800s, what technology has challenged over the past decades is the traditional viewpoint that robots are mere tools, rather than proper agents, in the legal field. Some reckon that we should register such machines just like corporations. This idea, for example, has been proposed by Curtis Karnow in *Liability for Distributed Artificial Intelligence* (1996), Jean-François Lerouge in *The Use of Electronic Agents* (2000) and Emily Weitzenboeck in *Electronic Agents and the Formation of Contracts* (2001). Certain scholars, as Anthony Bellia in *Contracting with Electronic Agents* (2001), suggest that we should bestow robots with capital. Others, as Giovanni Sartor in *Cognitive Automata and the Law* (2009), think that making the financial position of such machines transparent is a priority. Whilst further policies are feasible and even indispensable, *e.g.*, insurance models, what these proposals have in common has a precedent in the ancient Roman legal mechanism of *peculium*. According to the Digest of Justinian, the *peculium* was "the sum of money or property granted by the head of the household to a slave or son-in-power. Although considered for certain purposes as a separate unit and so allowing a business run by slaves to be used almost as a limited company, it remained technically the property of the head of the household" (Watson 1988: xxxv–xxxvi).

As a sort of proto-limited liability company, the *peculium* aimed to strike a balance between the claim of the masters not to be dilapidated by their slaves' businesses and commercial activities and the interest of the slaves' counterparties to safely transact with them. Most of the time, a master's liability was limited to the value of their slave's *peculium* and yet, the legal security of the latter guaranteed that obligations would have been met. For example, the contractual counterparties of the slaves could check whether the negotiations fell outside the authority or financial autonomy of the slave and, *vice versa*, in the wording of the Digest, "anyone who does not wish contracts to be made with him may prohibit it" by giving public notice (*Dig.* XIV, 3, 11, 3). Similarly, the mechanism applied when "the party desired business to be transacted with him under a certain condition, or through the intervention of a certain person, or under a pledge" (*Dig.* XIV, 3, 11, 5). But, going back to the case of the *institor* managing different classes of convenience stores, what did giving public notice mean?

> To give public notice we understand to mean that it shall be made in plain letters, so as to be easily read from the ground; that is to say, in front of the shop or place where the business is carried on, not in a retired place, but in one which is conspicuous. Shall the notice be in Greek or in Latin letters? I am of the opinion that this depends upon the character of the place, so that no one can plead ignorance of the letters…
>
> It is essential that the notice should be permanently posted; for if the contract was made before the notice was set up, or it was concealed, the Institorian Action will be available. Hence, if the owner of merchandise posted a notice, but someone removed it, or through age, rain, or something of this kind, the result was that there was no notice, or it did not appear; it must be said that the party who made the appointment will be liable. If, however, the agent himself removed it for the

> purpose of deceiving me, his malicious act should prejudice the party who appointed him, unless he who made the contract also participated in the fraud (*Dig.* XIV, 3, 11, 3–4. Trans. by S. P. Scott, *The Civil Law*, IV, Cincinnati, 1932).

Matters of legal certainty, financial and contractual warranty, or transparency, can obviously be improved in the case of modern autonomous robots. When following the example of ancient Roman lawyers, however, we should distinguish different kinds of robo-traders, as Romans did with multiple types of activities and status of the slaves as *dispensatores*, *ordinarii*, etc., for each of which specific lawsuits or *actiones* were established: besides the aforementioned *Institorian* action, think about the *actio exercitoria*, *tributaria*, etc.[3] Therefore, we have to distinguish the kind of business or commercial activity the robot is entitled to pursue, whether the robot acts on its masters' behalf or as a mediator between third parties, while being understood that the behaviour of the robot will be bound by rules and conventions that usually apply to the circumstances of the case. Consider the (not too futuristic) case of a robotic personal assistant such as a sort of i-Jeeves that helps us schedule a set of conferences, lectures and meetings at several European (or US) universities. Whereas we may guess at the best way of accepting simultaneous invitations from Oxford, Barcelona, Heidelberg, Athens and Paris, our robot needs not resolve the travelling professor problem by determining the shortest possible tour that visits each university only once. Rather, we expect that i-Jeeves checks both the availability and convenience of logistics in accordance with a number of parameters such as budget, time efficiency, or weather average conditions: i-Jeeves reports its findings back for a decision or, even, could determine the steps of our tour by directly booking hotel rooms, flights and so forth. Such contracts would not only be valid but, thanks to the digital *peculium*, a fair balance would be struck between the different human interests involved. By employing robots or artificial agents to do business, transactions or contracts, individuals could claim a liability limited to the value of their robots' portfolio (plus, eventually, forms of compulsory insurance), while the robots' *peculium* would guarantee their human counterparties, or other robots, that obligations would really be met.

On the other hand, we can further the Roman legal framework by granting robots legal accountability. As occurs with traditional artificial persons, as seen above in Sect. 2.3.2, legal systems may sever the responsibility of designers, manufacturers, operators and users of robots dealing with third parties, so that, on the basis of the warranty of their own *peculium*, only robots would be held liable for damages caused by them. Admittedly, this solution has several advantages: on the side of the contractual counterparties

[3] For a more complete list, see Ŝtaerman and Trofimova (1975: 82).

of robots, the personal accountability of such machines renders irrelevant whether they are acting beyond certain legal powers and who should be held liable for conferring such legal powers. On the side of users and operators, the personal accountability of robots allows humans to evade responsibility for possible malfunctions of the machine as well as errors of induction and specification, as seen above in Sect. 3.3.2. Moreover, aside from the quantification of the *peculium* and data on which insurance policies might hinge, the personal accountability of robots seems to be particularly recommended for certain applications. In light of a new generation of AI chauffeurs and intelligent car sharing, let me examine this hypothetical separately in the last section of this chapter.

4.5 The UV Revolution

One of the most dynamic fields of robotics technology today deals with the design, production and use of Unmanned Vehicles ("UV"). Although the technology is currently more prominent in the military than the civilian sector, a number of factors such as inter-agency transfers, increasing international demand, public R&D support and growing access to powerful software and hardware, explain why the civilian use of this technology is rapidly and progressively mounting. This is the case for several UV applications such as for border security, law enforcement, emergency and hazard management, remote exploration works and repair, urban transport, farming and more. As Brendan Gogarty and Meredith Hagger argue in *The Laws of Man over Vehicles Unmanned* (2008), the relative cost savings promised by UV technology have "excited many commercial operators" (*op. cit.*, 110), so that it is crucial for lawyers to assess the regulatory constraints for the ever-growing production and use of this new generation of UVs. More particularly, attention should be paid to three types of unmanned vehicles.[4]

[4]As mentioned in Sect. 3.5, we should grasp the unmanned vehicles as part of a more complex multi-agent system where such autonomous or semi-autonomous machines interact with maintenance and safety contractors, traffic operators or internet controllers, in order to avoid communication interferences, environment concerns, collisions, and the like. By considering that such machines will increasingly be connected to a networked repository on the internet that allows robots to share the information required for object recognition, navigation and task completion in the real world, some scholars refer to this type of robots as intelligent unmanned systems, unmanned aircraft or rotorcraft systems, and so forth. The aim of this section, however, is to stress the different ways UAVs, UUVs, and UGVs may affect current legal frameworks, rather than the systemic features of such network-centric applications.

The first type is provided by aerial applications, that is, UAVs. As previously stated above in Sect. 3.3, more than forty countries are currently developing such a kind of technology for military purposes. In addition, there already are cases of non-lethal engagement of suspects, arrests by drones, monitoring operations and UAVs specifically designed for policing, patrolling and inspection. As Peter Singer stresses in *A World of Killer Apps* (2011), "police departments in cities such as Miami, Florida and Ogden, Utah, have sought special licenses to operate unmanned aerial surveillance systems." However, the advancement is so rapid that drones already are within the reach of public bodies, private companies and even individuals. Both the US and EU are adopting regulations and procedures so as to permit UAVs to share the same airspace as commercial traffic. Aside from the law enforcement field, consider the definition of aircraft and related products as contained in Article 3 of the European Regulation EC 216/08, which appears broad enough to include UAVs. Likewise, in the spring of 2011, the US Congress established that "US civilian airspace should be opened to allow more widespread use of such systems by 2015" (Singer 2011). Rather than issues of military immunity and criminal accountability as previously mentioned in Sect. 3.5, the civilian use of UV technology puts forward problems of human responsibility and contractual liability concerning safety claims such as control loss, link issues, automated recovery or piloting regulation.

The second type of UV technology is offered by water-surface and underwater ("UUV") applications such as in remote exploration work and repairs of pipelines, oil rigs and so on. Among UV devices, this is one of the most developed fields: Gogarty and Hagger have even spoken of the golden age of UUV technology that "occurred more than a decade before the UAV revolution" (*op. cit.*, 104). Whilst development in UUVs and the increase of their use in the civil sector are likely to force lawmakers to amend many clauses of the current legal framework in maritime law, *e.g.*, the 1972 IMO COLREGs Convention, it nonetheless seems that UUVs do not really affect basic tenets of the law. In light of today's spectrum of robotics applications, as seen above in Sect. 4.1, UUVs are in fact closer to reasonable safety and controllable machines such as the da Vinci surgery system, than the ultra-hazardous activity of (certain types of) UAVs. Although there are UUVs that autonomously undertake their work by preventing damage, alerting controllers or repairing oil rigs in the Caribbean Sea, the legitimacy of such automatic devices can be grasped by lawyers using the same concepts developed for previous technological innovations, that is, in terms of the probability of events and the cost of their consequences.

The third type of UVs finally offers some of the most challenging applications of this technology, namely, the civilian (rather than military) use of unmanned ground vehicles. Whether or not future UGVs will need driving licenses, special licenses, etc., UV cars and AI chauffeurs allow us to deepen the legal issues that are raised by the civilian use of both UAVs and UUVs. The complexity of the environment that designers and producers have to address increases the uncertainty and unpredictability of UGVs automatically driving on the freeways. As a matter of risk, these UVs are more similar to unmanned flying vehicles than unmanned ships exploring the deep ocean floor. Yet, contrary to the use of UAVs patrolling the air for law enforcement purposes, the risks of employing UV cars mostly regard contractual obligations and problems related to strict liability in the field of torts, rather than constitutional safeguards and human rights law. On this basis, proponents of UGV technology ask for “a major review and clarification of existing civilian traffic safety regimes and even the creation of a specific regulatory system for UVs” (Gogarty and Hagger 2008: 121).

The next section dwells on whether new forms of accountability for the behaviour of these machines, such as the digital *peculium*, fit the new generation of AI chauffeurs and intelligent cars. Then, in the final Sect. 4.5.2 of this Chapter, the focus is on how UGVs suggest that lawyers will increasingly address (or be pressed by) cases of extra-contractual responsibility, *e.g.*, robots damaging third parties rather than affecting their contractual counterparties. This scenario proposes a further type of responsibility, such as the *Aquilian* protection in Roman law.

4.5.1 AI Chauffeurs and Intelligent Car Sharing

Intelligent vehicles driving themselves on highways are a popular subject of Sci-Fi movies: over the past 50 years, however, a number of states, organizations and private companies have made the dream come true. In the 1960s, the idea of building fully autonomous UGVs has been seriously pursued in several countries such as the US, Japan, Germany and Italy. Two decades later, the European Commission began funding a project on autonomous vehicles, the Eureka Prometheus Project (1987–1995). In the late 1990s, the US Congress authorized the Defence Advanced Research Projects Agency (“DARPA”) to organize a series of prize competitions for driverless cars in order to develop the military sector of UGVs and make one-third of ground military forces autonomous by 2015. Whereas there already is a panoply of US military UGVs such as TALON and Panther M-60 (see Singer 2009), the advancement of the civilian sector has been impressive.

Consider the aforementioned DARPA Grand Challenge competition. The first race was held on 13 March 2004, in the Mojave Desert, but none of the cars completed it. Just a year and a one-half later, five vehicles successfully finished the second race. Starting a rivalry such as the competition between Oxford and Cambridge in the annual boat race, the 2004 winner, *i.e.*, the Carnegie Mellon University's Red Team was defeated by the Stanford University's Racing Team on 8 October 2005. Two years later, Carnegie Mellon had the opportunity to take the revenge at the "Urban Challenge." On 3 November 2007, the third DARPA competition concerned a 96 km urban area race, to be completed in accordance with all traffic regulations and within 6 h. Due to the rapid advancement of technology, the challenge was not only to complete such a tortuous route, but to complete it as soon as possible. Teaming with General Motors in the Tartan Racing, Carnegie Mellon overtook the Stanford-Volkswagen car, taking 4 h 10 min and 20 s, at 22.53 km per hour, to cross the finish line first…

Three years later, in 2010, the European Commission promoted the "Intelligent Car initiative." As the corresponding website is keen to inform, the aim is to "imagine a world where cars don't crash, where congestion is drastically reduced and where your car is energy efficient and pollutes less." There are around 1.3 million mishaps and 41,000 people who die in car accidents on EU roads each year (whereas, in the US, more than 37,000 fatalities occurred in 2008). Besides, traffic jams impact on 10 % of the European major road networks and costs are estimated 50 billion per year, that is 0.5 % of EU GDP. Moreover, road transport accounts for more than one-quarter of the EU's total energy consumption. Therefore, in the phrasing of the Commission, "the Intelligent Car initiative is an attempt to move towards a new paradigm, one where cars don't crash anymore and traffic congestion is drastically reduced. Part of the i2010 strategy to boost Europe's digital economy, the Intelligent Car initiative is an answer to the need of citizens, industry and the Member States to find common European solutions and to improve the take-up of intelligent systems based on information and communication technologies ("ICT")."

Meanwhile, under the supervision of Sebastian Thrun, the director of the Stanford AI Laboratory and team chief of the robotic vehicle Stanley – which won the 2005 DARPA competition mentioned above – Google has been developing and testing its own driverless cars. As of 2010, such vehicles have driven 230,000 km with some human intervention and 1,600 km completely alone. A year later, lobbied by Google, the Nevada Governor signed into law a bill that, for the first time ever, authorizes the use of autonomous vehicles on public roads. Approved by the Nevada Assembly (36–6) and the Senate (20–1), the law amends certain provisions governing transportation and, furthermore, establishes that the Nevada Department of Motor Vehicles "shall adopt regulations authorizing the operation of autonomous vehicles on highways within the State of Nevada" (AB 511, June 2011). Although such regulations on safety and performances standards may take a long time, what is at stake here concerns experimental cars where "a human driver can override any error," as John Markoff reports in The New York Times, quoting some Google researchers.[5]

Still, it is a short step to envisage fully autonomous UGVs driving themselves in Nevada and, for that matter, spreading ubiquitously on public roads. However, despite rapid advancement of technology in key components of such cars as adaptive headlamps and cruise control, blind spot monitoring and driver checking systems, traffic sign recognition, pre-crash schemes and so forth, it is likely that lawyers should be prepared to address a new class of hard cases. In fact, who should be liable if the autonomous car has an accident? In the phrasing of *The Laws of Man over Vehicles Unmanned*, how will fault be determined when a human and computer are sharing the reigns of a vehicle under traffic legislation? Indeed, who will be at fault if the vehicle has an accident when it is clear only the computer AI was in control? (Gogarty and Hagger 2008: 120–121). Moreover, in the name of urban sustainability and green policies stressed above, how about new forms of distributed responsibility as soon as we reflect on, say, schemes of AI car sharing?

As mentioned above in Sect. 4.3.2, traditional forms of apportioning individual liability fall short in coping with such scenarios. Let me insist on three points:

First, there is the difficulty for traditional legal outlooks of addressing the behaviour of robots as agents, rather than simple instruments of human interaction. As a matter of fact, humans will delegate to such autonomous and even intelligent cars complex cognitive tasks, such as driving themselves on the highways, while avoiding other cars, preventing individuals' reckless manoeuvres and so forth.

Second, from the fact that a human let the car drive by itself, it does not follow that the legal effects of the decisions of that car should necessarily fall upon the human. On the one hand, we are back to cases of apportioned responsibility of designers, manufacturers, dealers and users of AI machines, which inspired Curtis Karnow to predict a failure of legal causation as discussed above in Sect. 3.5. On the other hand, the hypothetical of environmentally friendly-AI car sharing makes this scenario still more complex, since such machines would be dealing with a multitude of human masters.

Finally, we should take into account the protection of third parties. Compared to the form of agency in the case of robo-traders, the spectrum of third parties widens so as to transcend the field of contractual obligations and concern what common lawyers call torts, that is, in the jargon of civil

[5] *Google Cars Drive Themselves, in Traffic*, October 10, 2010, A1 of the New York edition.

lawyers, forms of extra-contractual liability. In the case of robo-traders, individuals grant them authority to act on their behalf when dealing with third parties, so as to accept bids, make offers, compare prices, etc. In the case of AI chauffeurs, individuals will grant them authority to autonomously drive on the freeways, so that, theoretically speaking, everybody could be affected by the reckless behaviour of these machines.

A new form of accountability, such as the digital *peculium*, that could successfully tackle the legal challenges of a new generation of UGVs was introduced in Sect. 4.4.1. After all, we can imagine AI chauffeurs that accept offers, or make contracts, so as to autonomously drive individuals on the streets. Therefore, on the side of the contractual counterparties of robots, the personal accountability of AI chauffeurs guarantee that obligations for damages caused by such machines would be met. On the side of both users and operators, the personal accountability of AI chauffeurs let people evade liability for possible unpredictable malfunctions of the machine. Whilst it is crucial to determine the sum of money granted to the intelligent car, it is likely that programs such as Google's driverless cars or the European Commission's i2010 strategy will provide enough data on the probability of events, their consequences and costs, to determine levels of risk and, therefore, both the amount of the *peculium* and forms of compulsory insurance, on which new forms of accountability for the behaviour of such machines may hinge. This is the approach suggested by a number of scholars, such as Tom Allen and Robin Widdison in *Can Computers Make Contracts?* (1996), Ian Kerr in *Ensuring the Success of Contract Formation in Agent-Mediated Electronic Commerce* (2001), Woodrow Barfield in *Issues of Law for Software Agents* (2005), Francisco Andrade et al. in *Contracting Agents: Legal Personality and Representation* (2007), down to the aforementioned works of Giovanni Sartor (2009) and Chopra and White (2011).

However, would new forms of personal accountability for robots represent the one-size-fits-all answer to the new generation of legal issues brought on by such robots? Does this approach apply equally to robots as agents and robots as instruments? Does the legal accountability of the robot suffice to deal with different types of claims in the field of torts?

4.5.2 Unjust Damages

We have examined three different types of robots in this Chapter. First, we dwelt on robots as means of human industry and interaction that include both ends of the spectrum of robotic applications as examined in Sect. 4.1; namely, reasonable safe and controllable machines, such as the da Vinci surgery system, and the ultra-hazardous activities performed through some of today's UAVs. As means of human industry, such machines do not challenge basic tenets of the law as current provisions of contracts and tort law properly address damages or harm caused by these robots. Think of strict product and malfunction liability claims, breach of warranty, negligence, or evidence, that is, the set of concepts examined through the mechanism of the burden of proofs in the *Mracek v. Bryn Mawr Hospital* case discussed above in Sect. 4.2.2. As Richard Posner affirms in *Economic Analysis of Law* (1973), "new activities tend to be dangerous because there is little experience with coping with whatever dangers they present… The fact that the activities are new implies that there are good substitutes for them" (*op. cit.*, 2007 edition: 180).

A second class of robotics applications has to do with robots as legal agents. Rather than simple objects concerning clauses and conditions of contracts, the example of certain robo-traders has shown machines capable of determining clauses and conditions of contracts by themselves. Here, current provisions of the civil (as opposed to the criminal) law fall short in addressing both the cognitive states of such machines and ways for determining or apportioning liability for damages caused by this class of robots. Some ways for severing the chain of responsibilities between designers, manufacturers, operators, users and third parties that interact with such machines, were discussed above in Sect. 4.2.2 and Table 4.1, according to three different kinds of erratic behaviour: robotic specification, induction, and malfunction of the robot. Whereas traditional legal standpoints end up in a Hegelian night, where all kinds of liability are blurred into the same grey colouring, we should define where to cut back on the scale of the activity. New forms of accountability for robots as strict agents in the civil law field, *e.g.*, the digital *peculium*, show how to prevent this threat, so as to "cope with whatever dangers they present" (Posner 2007). By granting authority to the robot, so as to let it act on an individual's behalf when dealing with third parties, a new form of *peculium* strikes a fair balance between the counterparties of robots demanding the ability to safely interact or transact with such machines and individuals claiming that they should not be ruined by the decisions or behaviour of their own robots. Although it would be meaningless to treat the first class of robots, *i.e.*, robots as means as legal persons with a contracting capability in their own right, it makes a lot of sense to attribute such capability to the new generation of robo-traders.

Finally, there is the class of robots as intermediates in social life, rather than agents of human business and negotiations. As the example of the AI chauffeurs has shown, such robots can make business and still most of the time, they will be dealing with third parties, namely, individuals who are not directly concerned by the enforcement of rights and obligations created by the robots' business. In the phrasing of the UN 2005 Robotics Report, this class of machines concerns "domestic or personal use of service robots for domestic tasks, entertainment, handicap assistance, personal transportation, home security and surveillance." Such a class of robots as intermediates of human interaction brings us back to the scenario of AI chauffeurs provoking accidents on the highway. Consider a new generation of robot toys (entertainment), or robot nannies (domestic tasks and handicap assistance). In the case, say, a nanny such as Jetsons' Rosey, nursing your old mother, causes harm to some of your mother's acquaintances, who is liable?

This scenario goes beyond the contractual mechanism of *peculium* and involves what Roman jurists defined in terms of *Aquilian* protection; namely, the form of responsibility stemming from the general idea that individuals are held liable for unlawful or accidental damages caused to others because of their personal fault: *Alterum non laedere* as discussed above in Sect. 2.2. Although the digital *peculium* may govern certain cases of extra-contractual responsibility, *e.g.*, road accidents, there is a number of further obligations, so as to protect from unjust damage, in the many-to-many, rather than one-to-one contractual scenarios of social interaction. Think of strict liability rules in the field of robotics by analogy with dangerous animals as seen above in Sect. 3.4.3. Likewise, consider cases of liability for the negligent control of artificial agents and even vicarious responsibility for the autonomous acts of individuals' artificial employees. What is crucial here concerns the different robotic applications with which we are dealing, since such robots as domestic service robots, as a sort of AI children, animals, or i-Jeeves, entail different types of liability and opposite ways to determine on whom the burden of proof should fall. These are cases where we need a further type of expertise in the laws of robots. After the chapters on crimes and contracts, we will deepen the examination of that which common lawyers define as the field of torts.