On automatic causal reasoning for legal analysis

Jos Lehmann and Joost Breuker
Department of Computer Science and Law
Faculty of Law
University of Amsterdam, The netherlands
{jos, breuker}@lri.jur.uva.nl

Abstract. The present article offers an overview of the set-up of our research project about the automation of causal reasoning for legal analysis. Firstly, we formulate an argument against the undermined role of causation in the view of some legal theoretical schools. Secondly, we present the most relevant approaches to the problem and we justify our choice of Hart and Honoré's account as a jurisprudential guideline throughout our research. Thirdly, we select Pearl's inferential framework as a computational support for counterfactual reasoning. Finally, we point out the problems we are currently dealing with, in an attempt at merging Hart and Honoré's qualitative account of causation and Pearl's formal theory of causality.

1. Introduction

In the framework of our research project, which investigates the "Recognition of intention and of responsibility in the semiautomatic analysis of a legal case", we have been treating various and disparate aspects of the topic. On one hand, we deal with what can be called "external" or methodological questions. As a matter of fact, we assume the legal theoretical position that sees the attribution of responsibility as a form of inference based on common sense causal reasoning. Not everybody shares our sincere convictions, though; which has forced us to study them better and to find support in the literature. We, hence, propose such literature both as the methodological justification of our research and as a common ground of confrontation and of further acquisition of relevant scientific material. On the other hand, we have to face "internal" or ontological problems, in the attempt to define the minimal conceptual apparatus required by any plausible form of (automatic) recognition of responsibility. The search for this minimal set of concepts is performed following three steps:

1. *individuation* of the common sense and of the legal-theoretical concepts that are usually employed by legal experts (e.g. judges, lawyers, etc.) in reasoning about the attribution of legal responsibility to the agents involved in a case;
2. *formal definition* of these concepts;
3. *automation* of their normal usage by means of dedicated inference mechanisms.

The external and internal research activities mentioned above are all directed at the single aim of expanding the Functional Ontology of Law (FOLaw in the following) [11], a core ontology of legal knowledge, developed as a support to the engineering of legal expert systems.

Jos Lehmann and Joost Breuker, 'On automatic causal reasoning', in: Joost Breuker, Ronald Leenes and Radboud Winkels (eds.), *Legal Knowledge and Information Systems. Jurix 2000: The Thirteenth Annual Conference.* Amsterdam: IOS Press, 2000, pp. 123-134.

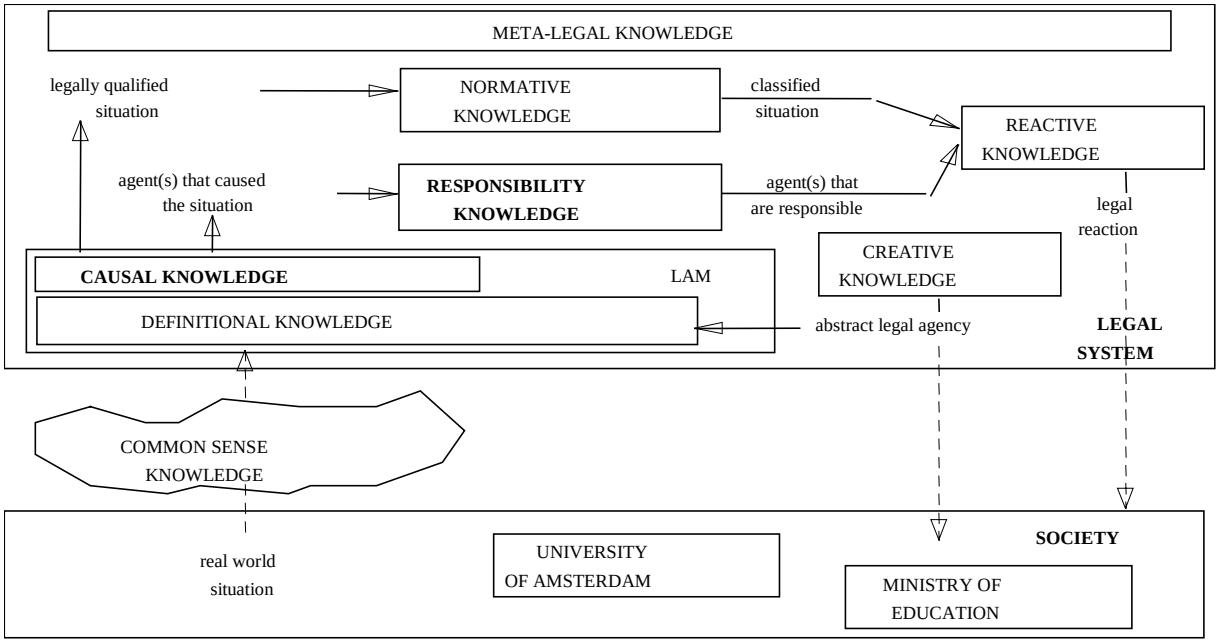

Figure 1: The Functional Ontology of Law

FOLaw tries to capture and to formalize both the internal structure and the functioning of a legal system. It relies on two main (modeling) assumptions:

- Legal knowledge can be "modularly" represented, i.e. a legislation consists of different categories of knowledge which can be distinguished from each other;
- Legal knowledge can be functionally represented, i.e. the different modules fulfil different functions both within the legal system and in the analysis of a single case.

Figure 1 (from [11], p.74) shows the different categories of legal knowledge that Valente has "isolated" and their functional dependencies. These dependencies form the typical structure of legal argument, which can, for instance, be executed as a series of steps in assessing a legal case (but the same structure may also support the drafting of legislation).

Our main objectives have so far been the *individuation* and the *definition* of causal knowledge. We follow a rather broad and qualitative method of research: conceptual analysis. This is due to three main factors: the *nature of the treated subject* (i.e., legal causal reasoning), which is an argumentative kind of activity, definitely based on a qualitative image of the world; the need of *expanding an ontology*, which is defined as the specification of a conceptualization; and the *seek for generality*.

The present article offers an overview of the set-up of our research project. Section 2 contains a counter-argument against the undermined role of causation in legal theory. Furthermore it gives an overview of the most relevant legal theoretical approaches to the problem. Based on such overview we choose Hart and Honoré's account [5] as our jurisprudential guideline. Section 3 briefly goes into the fields of philosophy and artificial intelligence, where we actually stayed a long time in search of inspiring theories. One of these is Pearl's inferential framework, which we adopt as computational support for counterfactual reasoning. Finally, Section 4 shows the problems we are currently dealing with, in the attempt at merging Hart and Honoré's qualitative account of causation and Pearl's formal theory of causality.

2. Causation in Legal Theory

2.1 Dubitant ergo est[1]: a counter-argument against external criticisms

The problem of formally, consistently and completely defining the relation of causality between two events (or states of affairs) is one of the hardest among the scientific and philosophical questions of all times. Every kind of formal theory of the world (e.g., physics, chemistry or economics) and any organized corpus of knowledge about the world (e.g., the law, medicine or history) suffers from internal paradoxes which are provably due to some lack of understanding of causal relations.

The causal plague is so diffused and general that each scientific field has developed its own particular rules of the thumb for dealing with causal issues. As a matter of fact, specific causal rules have so far proven to be much more reliable than general inferential schemes. The latter should in principle guarantee more uniformity in distinguishing causal from non causal relations; but the former actually provide a tighter correspondence between general causal intuitions and each field's peculiar way of describing basic facts, of assessing their truth and (most importantly) of handling exceptions. This sort of "causal scientific secessionism" has generated a big deal of technical views on causality which are hardly comparable to each other. Nonetheless, as the word *cause* and its synonyms keep being used for communication within and across different fields, secessionism has also produced a general desire for unity in the form of "federalism"! Every (sufficiently theoretical) study of causality almost invariably makes two statements: first, the word *cause* refers to a cluster of variegate and sometimes even contradicting concepts (rather than just one clear-cut notion); second, the unity of such diverse conceptual family is guaranteed by a common sense lattice, i.e. by something that everybody has (i.e., common) but that nobody can rationally describe (i.e., sense).

The law makes no exception, neither for its (sometimes provincial) tendency to secessionism nor for its desire to fully take part to the causal federation. Going through (even only a part of) the legal-theoretical debate on the problem of causation makes very clear how deep is the dichotomy between different groups of legal theoreticians. Some, the secessionists, sincerely think that legal practice and reasoning has no problem whatsoever with causation. In their eyes, everything in the law has only to do with attributing legal responsibility; in other words, legal reasoning deals only with so called questions of policy. On the other hand, the federalists try to convince their fellow scholars of the following: if it is true that legal practice has felt the need to let juries or popular judges solve difficult questions of fact arising in legal cases; then it is also true that legal reasoning has problems with causality as much as any other type of reasoning has. Therefore, the pragmatic solution (i.e., letting juries decide questions of fact) should not be considered as a miraculous way out of the problems of causal reasoning. Juries do not shield the court room from the causal federation. On the contrary, they work as the channel of communication between the legal apparatus and social common sense, in any given moment of time. The function of juries is twofold: they represent in the court room *general* social knowledge (at least in the form doubts, if not as certainties) about questions of fact and they provide a reasonable solution for the *specific* case they are called to judge. In short, juries *at least* doubt; therefore, the problem, that they are called to doubt about, exists.

---

[1] They doubt, therefore it exists.

2.2 Problem definition and examples

If the problem of legal causation exists, then we would like to give a proper definition of it. We think that the analysis conducted in [5] is the best starting point on the subject. The book has enough attention for theoretical studies on causality and enough interest in real-life legal cases for attempting to bridge the gap between "pure causal reason" and "practical legal causal reason".

The two English authors provide the following general definition of the problem of causal reasoning:

Definition 1: How is it possible to *uniformly* and *consistently* establish *legal responsibility* by means of *a standard form* of *causal reasoning* applicable to the widest possible set of different cases?

The definition given above can be specialized for tort, contract and criminal law. This is done just by inserting in the general definition the specific terminology of each of those branches, because "causal limitations are to a great degree the same in tort, contract and criminal law."(in [5], p.132). Given those similarities, in the rest of this paper Definition 1 is adopted as the working problem definition.

Besides the problem definition, we present here a case which constitutes a severe challenge for any theory of causal reasoning. Tracing the chain of causation for this story and making a decision about the attribution of responsibility is a major problem, even if the decision is purely based on questions of policy. The purpose of introducing this case here is to stimulate the reader's imagination with a "causal monster". Further in the paper, the Case will be used for illustrating and testing various theories of causality.

Case for tort law: "In breach of a statute forbidding the sale to an infant under the age of 16 of dangerous weapons, the defendant sold an air rifle and ammunition to a boy of 13. The boy's mother told the boy to return the weapon to the defendant and get a refund: on the defendant's refusal to take the rifle back, the boy's mother took it from the boy and hid it. Six months later the boy found it and allowed a playmate to use it, who shot and accidentally wounded the plaintiff, destroying the sight of one eye."[2]

As mentioned above, legal-theoretician have been having a hard time trying to find a principled way of reconstructing the chain of causation in the analysis of legal cases such as this (or even simpler ones). In the following section various theories are presented, which are particularly interesting for the purpose of this research. Their (increasing) level of formalization gives a good idea of how far is the automation of legal analysis from the present stage of scientific and technological development.

2.3 Main legal-theoretical approaches to the problem of causation

We have already briefly introduced the distinction between two main legal-theoretical positions on the problem of causation which we metaphorically referred to as federalism and secessionism. In legal theory they are usually known as the *traditional view* (federalism) and *causal minimalism* (secessionism). The traditional view adopts the principle of *causal proximity* as the guideline in the causal analysis of legal cases. According to this view, *the* cause of a harm is the most proximate among all the conditions that made it possible for the harm to come about. Proximity is not only meant as a temporal distance, but also as a

[2] Henningsen v. Markovitz (1928) 132 Misc. 547, 230 NYS 313

material or factual one. The judging authority should therefore establish such proximity based on a common sense reconstruction of the case. Difficult examples can further be handled by dedicated "rules for the determination of proximate cause"(in [5], p.96) such as those proposed by Beale in the beginning of the 20s. In the 30s a new wave of legal theoreticians (e.g. Green or Gregory) opposed the very possibility of distinguishing between questions of fact and questions of policy. They claimed that most causal issues in legal analysis are solved by a combination of legally defined concepts (e.g., risk) and *sine-qua-non* (i.e., counterfactual reasoning). Only few hard cases should be handle by common sense causal reasoning.

Due to the publication of Hart and Honoré's book, the end of the 50s saw a renewed interest in legal causation. The authors argue both against the traditional view and against causal minimalism and they propose an alternative view, which is a synthesis of the two opposed theories.

The arguments of the two scholars against the existing doctrines of legal-theoretical causation are somehow intertwined. On the one hand, the notion of causal proximity is too vague for supporting a coherent form of legal analysis. Therefore, the definition of a clearer concept of causation is needed. This should not become, though, an attempt at defining a set of rigid "rules for the determination of proximate cause", because most of these rules for handling questions of fact implicitly encode legal principles, i.e. answers to questions of policy. On the other hand, Hart and Honoré advocate that the correct usage of technical legal concepts such as *liability*, *foreseeability*, *risk* and *scope of the rule* is highly dependent on (the application of) a plausible notion of causality; and this notion, they say, is mostly hidden by causal minimalist in the *sine-qua-non* test (i.e., counterfactual reasoning), which is proposed as a necessary inferential counterpart of the technical notions mentioned above.

In order to ease the tension between the common sense perception of the world and the (technical) legal view on it, Hart and Honoré propose a whole new approach to the problem. The main methodological novelty of their study is the idea of importing analytical philosophy into the discussion on legal causal reasoning.

The starting point of their analysis is the observation that the lack of agreement between the traditional approach and causal minimalism is due to a misunderstanding about which part of the so called *judicial viewpoint* should be given more consideration in legal causal analysis. Such viewpoint assumes the following:

1. In a case-description the causal, factual links between events are considered as (intuitively) detectable;
2. Legal causal reasoning is concerned only with *agent causation*, i.e. with establishing which of the mentioned *agents* can be labeled as *the* responsible one for a specific event (i.e., the harm) reported in the case-description.

Now, while the traditional approach attempts at giving a clear encoding of the first assumption, causal minimalism concentrates on encoding the second assumption, dismissing the first as a pure matter of intuition, hardly interesting for "real-life" legal analysis.

The legal language of causation must hence be studied by means of discourse analysis, which, at the time when Hart and Honoré wrote, was a major development in the Anglo-Saxon philosophical landscape. Such study should refrain from both temptations of explaining everything (as in the traditional view) and of leaving everything unexplained (as in causal minimalism). Furthermore, it should avoid another traditional philosophical mistake: reducing causal reasoning to its scientific form, i.e. induction. Legal analysis is a qualitative kind of reasoning, which shares more objectives with history than with exact sciences. A legal expert looks for *an* explanation of one *particular* event, never for *the* general law

explaining *classes* of events! Hence, Hart and Honoré propose a causal inference scheme, based on a wide analysis of legal causal language in a (non arbitrary) selection of cases.

Table 1: The identification of events for Case 1

Event's Name	Process	Object	Time	Actor
e1	selling	rifle	t1	defendant (D)
e2	ordering	boy	t2	mother (M)
e3	refusing	boy	t3	Defendant
e4	hiding	rifle	t4	Mother
e5	finding	rifle	t5	boy (B)
e6	allowing	mate	t6	Boy
e7	shooting	bullet	t7	playmate (P)
e8	wounding	plaintiff	t8	

They see their proposal as a guideline for an aware use of causal reasoning by the court; they think of it as a moderate way of bridging the gap between the two extreme interpretations of the judicial viewpoint. The two English scholars provide the following definition.

Definition 2: Agent *A* causes an event *e*, that might involve agent *B***,** if either:

1. *A* starts some *physical process* that leads to *e*;
2. *A* provides *reasons* or draws attention to reasons which might influence the conduct of *B*, who causes *e*;
3. *A* provides *B* with *opportunities* for doing things that will lead *B* to cause *e*.

All three types of causation "have important negative variants."(in [5], p.3). For illustrative purposes Definition 2 is here applied to the Case. By doing this the virtues and the vices of Hart and Honoré's proposal are made clearer both at the modeling and at the inferential level. To model the Case, first, chop up the story into a temporal sequence of events as in Table 1.

The two English author do not provide an explicit specification of what they mean by event. Therefore for the moment we use a widely acceptable ontology of event based on objects and processes Second, "immerse" the events into the space of preconditions (i.e., a conceptual space where every event is a precondition of all the successive events represented by the box in Figure 2), ordering them by actor and time. Third, make an inferential step: for each event control if it can be considered as a cause of any successive event according to Hart and Honoré's test. Roughly speaking, such assessment consists in establishing whether the process of the precondition "leads" (physically of psychologically) to the object of the postcondition. In our example, the resulting graph might look as Figure 2.

In the intention of Hart and Honoré, a causal reconstruction in terms of physical processes, reasons and opportunities should clarify the material development of events and, therefore, help in applying legal principles. In other words, it could be a good way of focusing the attention of the legal expert on those parts of the case that are particularly relevant for determining responsibility. On the other hand, the model leaves a lot of unsolved ambiguities. These ambiguities are intentionally left unsolved by Hart and Honoré; because, as explained, the scheme they propose should only work as a guideline in legal analysis, not as a set of rigid rules directly applicable to every case. It is hence the practitioner's job to resolve the uncertainties in the meaning of the terms in the scheme, every time this is applied. But for the purpose of automation a more precise definition of the notions used in the scheme must be provided.

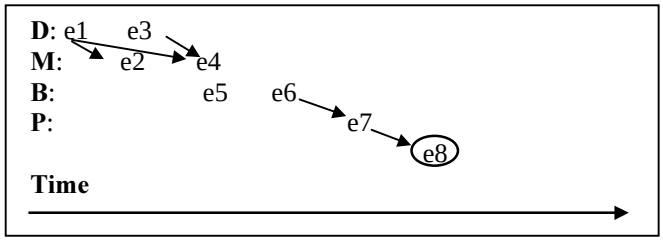

Figure 2: A model of Case 1 in Hart and Honoré's theory

Hart and Honoré's approach to legal-theoretical problems inspired other studies on the problem of legal causal analysis. The most recent result produced by this stream of legal-theoretical research is [2]. The Swedish authors propose a complex formalism for the establishment of liability in tort law. Their legal-theoretical position on legal causality can be labeled as minimalist.

Åqvist and Mullock's formalism is a blend of logical, game-theoretical and probabilistic techniques. Agent causation is determined by evaluating the *intent*[3] of the agent(s) while performing the action(s) which contributed to the occurrence of the wrong. The evaluation of the intent is accomplished by looking back in the causal chain ending with the wrong until a voluntary human action is found (first parameter: want) and analyzed with respect to the foreseeability of the wrong (second parameter: knowledge). The degree of foreseeability of the event is directly dependent on the *probability* of the course of events which the wrong belongs to. In Åqvist's system such probability factor is assigned to different courses of events at the beginning of the analysis of a case. It somehow represents the objective (in other words, common sense) knowledge used for the analysis itself.

For brevity we do not describe Åqvist's system in detail and we therefore avoid to model the Case. The main assets of Åqvist and Mullock's system are the high degree of formalization provided by their models, the large set of cases they use to test their theory, and the methodological example they establish in integrating legal-theoretical, philosophical and technical sources on the subject of causation. On the other hand, the fascinating simplicity of their system can easily turn simplistic, especially while modeling a case. Just as Hart and Honoré' s proposal, Åqvist and Mullock's leaves a lot of unsolved ambiguities, which are a major problem for their formalism and for any formal analytical tool used in detecting causal relations. Specifically, it is very legitimate to wonder how can we objectively and consistently decide about the foreseeability of courses of events. Åqvist and Mullock simply avoid this kind of questions by adopting a strictly minimalist interpretation of the judicial viewpoint.

2.4 Outside legal theory

Among the four presented approaches we see Hart and Honoré's as the most suited for our purpose. Building a (correct) formal theory of the qualitative knowledge present in Hart and Honoré's inference scheme would be a successful result. Such a formal theory would make the semi-automation of legal causal reasoning a feasible option. It could make it possible to build a software component that, by manipulating formulas that contain knowledge about physical processes, reasons and opportunities, causally analyzes case descriptions and assigns legal responsibility to agents. As a matter of fact, such a language would be an analytical representation of legal reasoning and not just a probabilistic simulation of it. Its design,

[3] Intent is the practical counterpart of intentionality and it consists of want and knowledge.

though, requires at least some knowledge of the existing philosophical and AI theories of causation. This is necessary in order to make the best choices we can, despite the present limitations in the representation of common sense. The next section presents a highly condensed overview of what philosophy and AI have to say on causation, with particular attention to the latest development in the latter field, i.e. Pearl's approach.

3. Causation in Philosophy and in Artificial Intelligence

As far as philosophy is concerned, Section 2 on legal theory gives enough of an idea of what a philosophical dispute looks like: vague, prolix, metaphorical and confusing. If we had to try the same type of reconstruction for philosophical accounts of causal concepts, we would end up writing an article within the article. Instead we present here two lists: the list of problems and the list of the approaches. On the one hand, two thousands years of philosophy have produced a set of questions that any minimally complete theory of causality must answer. These questions concern: ontology of the *relata*; causal principles, laws and statements; reductionism; experience of causation; direction of causation; the role of agency. On the other hand, the same two thousands years have resulted in five distinct types of solutions: implicational; counterfactual; instrumental; probabilistic; singularistic. We hope that the terms listed here are suggestive enough of the profoundly different picture of the world that each theory offers. The interested reader can find more information on this topic in the literature, particularly in [6]. As a conclusive remark on philosophy, it must be noted that at least one form of agreement has been reached about causality: the terminological distinction between *causality* and *causation*. Almost every author agrees that it is necessary to distinguish between the *relation* that connects a cause to its effect, which is called causality, and the concrete *process* that brings about the effect from the cause, which is called causation. In this article we do not adopt the distinction; just as in legal theory, we *only* speak of *causation*. Therefore, when the term causality is used, it is only for aesthetic purposes, and *not* to mark any distinctions from causation.

As far as AI is concerned the list of problems contains four main items: prediction, explanation, diagnosis and planning. On the other hand, the list of approaches contains two families: the symbolic approaches (situation calculus, naïve or qualitative physics, non standard logics) and the numeric approaches (Bayesian networks). More information on this topic can be found in the references.

3.1 Pearl's new approach

What needs to be explained in more detail, is Pearl's theory of the *Actual Cause* [9], which will become the inferential framework of our language.

Pearl's most recent theory of causality is based on an enhancement of the deterministic "half" of Causal Bayesian Networks, i.e. causal models that could support a more constraint analysis of probabilistic information. Bayesian Networks are replaced by Functional Causal Models of the considered domain, i.e. by (the graphic representation of) sets of Structural Equations. These kind of equations were first introduced in the scientific practice of econometrics, genetics and social sciences. They combine deterministic and probabilistic information of the world, providing the necessary knowledge for counterfactual reasoning, which is an essential component of many forms of causal reasoning.
A structural equation has the following form[4]:

[4] In the following, capital letters are variables (e.g., X) while small letters (e.g., x) are their values.

$x_i = f_i(pa_i, u_i)$ $\quad\quad$ $i = 1,..., n$

where x_i is the value of any so called endogenous (i.e., observable) variable of the model, pa_i connotes the values of the set of variables judged to be the immediate causes of the variable X_i (the so called parents of X_i) and u_i are the values of the so called exogenous (i.e., unobservable) variables representing measurement errors due to omitted factors (i.e., non deterministic, probabilistic information). Finally, f_i are functions, which represent the structural (i.e., causal) development of the world.

Now, based on the notion of Functional Causal Model (the triple M = <V, U, F>) and on various other assumptions, Pearl has proposed many inference and modeling techniques. What all these inference mechanisms have in common is their reliance on the notion of *Structural Contingency*. Generally speaking, causal reasoning effectively takes place only if the structural knowledge of the model can partially be violated for the purpose of testing causal hypothesis. Therefore, in Pearl's formalism the *do* operator (which sets a variable to a chosen value) allows to "wipe out" equations from the model and to substitute them with some desired value. By means of this operation it is possible to test how the world reacts (or would have reacted) under some specific circumstances (other than the actual ones). This violation of the normal (or actual) course of events is called a Structural Contingency. The phrase, which almost sounds like an oxymoron, suggests two ideas: the (hypothetical) violation effects the profound laws of the model (therefore, it is structural); but it does it only for the temporary purpose of reasoning, of testing causal hypothesis (therefore, it is contingent). Among all the inference techniques based on structural contingencies, the notion of *Natural Causal Beam* is particularly important for the purpose of our research. A Natural Causal Beam is a refinement of a Causal Functional Model, where dependencies between variables are not "merely" functional. In a beam the parents of each variable's value are minimally specified, which means that each variable might have a larger or smaller set of parent-variables depending on the value we consider. This relation between specific values of families of variables is called Sustenance, and it is meant to tight up the logical notion of Sufficiency. Given a Natural Causal Beam of a domain or of a situation, it is sufficient to run a counterfactual test in order to detect causation in fact. Using the Case, we illustrate here how this technique works. Our starting point is, of course, a Functional Causal Model as specified in Table 2 an Figure 3. In our scenario both the exogenous, Ui, variables are set as follows: U1 = yes and U2 = no. The endogenous variables take the values that bring about the wounding. We can hence put the model the following counterfactual queries:

- *Ceteris paribus*, had the defendant not sold the weapon, would the plaintiff still be wounded?

Formally: $\text{Plaintiff}_{\text{Defendant} = \neg\text{selling}} = \neg\text{wounded}$

- *Ceteris paribus*, had the mother got rid of the weapon, would the plaintiff still be wounded?

Formally: $\text{Plaintiff}_{\text{Mother} = \neg\text{hiding}} = \neg\text{wounded}$

Table 2: Function specification for the Case

Function	Domain	Codomain	Condition	Value
-	-	Ui i = 1,2	-	yes no
f1	U1	Defendant	U1 = yes Otherwise	selling ¬selling
f2	Defendant ∪ U2	Mother	Defendant = selling ∧ U2 = yes Otherwise	hiding ¬hiding
f3	Defendant ∪ Mother	Boy	Defendant = selling ∧ Mother = ¬hiding Otherwise	finding ¬finding
f4	Boy	Playmate	Boy = finding Otherwise	shooting ¬shooting
f5	Playmate	Plaintiff	Playmate = shooting Otherwise	wounded ¬wounded

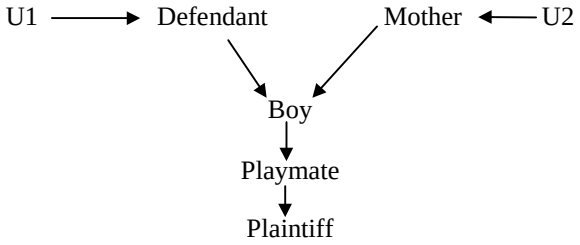

Figure 3: A Causal Functional Model of the Case

We get the same negative answer to both of them. This is correct, but insufficiently stringent for our needs, because we would like to go beyond the individuation of contributory cause and determine causation in fact, i.e. one chain of causes and effects. In order to do this we have to refine our model and transform it into a Natural Causal Beam. This is done separately for each possible scenario (i.e., for each possible value of the Ui variables). We have to determine the minimal set of parents that sustains (i.e., that is sufficient for) the actual value of a variable. In our model we just have one variable with multiple parents: Boy. We will hence focus on this one, as the actual cause of the other ones is the actual value of the parent variable. Now, in the real scenario we say that the variable Defendant sustains the actual value of Boy more than Mother does. This is because if we change the value of Defendant also the value of Boy must change; while setting Mother = ¬hiding does not invariably imply Boy = ¬finding. This consideration forces us to turn our model into the Natural Causal Beam shown in Figure 4 where f3 is modified as in Table 3, while all other functions remain unchanged. Running the counterfactual test on the beam gives us the actual cause of boy's finding the weapon: the defendant's act. This completes the chain of causation for the Case.

Table 3: New function specification for the Case

Function	Domain	Codomain	Condition	Value
f3 (modified)	Defendant	Boy	Defendant = selling Otherwise	finding ¬finding

Pearl's system (correctly) detects the actual cause. This, on one hand, is very encouraging: his inference mechanism works nicely. And not only this: in contrast with Åqvist and Mullock's formalism, Pearl's does not rely on a *possible world* semantics, which is, at least so far, highly demanding at the computational level. On the other hand, though, the positive

result should not hide the main difficulty of this as of many other systems dealing with causation: how to guarantee inferential consistency and uniformity over a larger set of different examples.

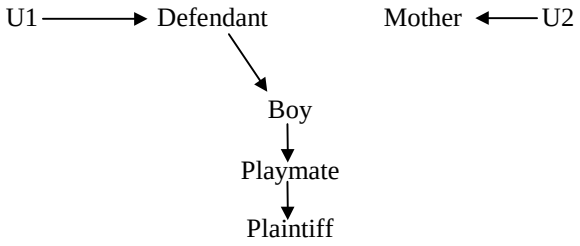

Figure 4: A Causal Functional Model of the Case

4. Conclusion

As we already pointed out for Hart and Honoré and for Åqvist and Mullock, also for Pearl two serious modeling questions arise:

- Do (sustenance-) functions between variables represent general knowledge about the world or specific knowledge about the considered case? They should encapsulate general knowledge about the world, but Pearl's modeling practice does not really deal with this issue;
- How can we compare the (relative) correctness of the output of counterfactual tests run on two different causal beams either for the same case or for two different cases? Is intuition enough of a criteria?

Even partially solving those two problems would be a step forward. The semantical reliability of the conclusions reached using Pearl's counterfactuals would increase. This means that, for the purpose of our research, Pearl's inference techniques has to be provided of a sort of modeling toolkit: an ontology of events and processes. This would specify:

1. The level of granularity of the knowledge which is allowed into the model;
2. General (sustenance-)functions for representing the admissible (material) changes of the described (physical) world, i.e. processes.

In other words, the proposed ontology will be a form of Naïve Physics, specialized for legal reasoning, which expands Pearl's ontology. The (minimal) set of concepts that our ontology will contain is the following:

- Event
- State of affairs
- Object
- Interval of Time
- Place
- Process
- Structure
- Start
- Finish
- Context
- Causation
- Relation
- Property
- Time
- Space
- Quality
- Compatibility

Furthermore, by looking ahead into our project, we see the attribution of responsibility defining itself as the activity of testing a causal reconstruction under a legal perspective. Legal reasoning goes beyond the question who caused what in a strict sense, and tries to identify a distribution of liability among the agents that are part of the causal structure that has lead to a violation of a norm. If the law does not state explicitly a default responsibility (e.g. a parent for a child) and there are not other violations of norms, the responsibility assignment follows the causal analysis, whereby the distribution of responsibility is (directly) proportional to the causal distance from the harm. However if another violation takes place as well (in our case the shopkeeper who sold the weapon to a minor and did not check his identity), there appears to be a serious aggravation or reinforcement of responsibility.

References

[1] Allen, J.F., 'Towards a general theory of action and time', in: *Artificial Intelligence*, Vol. 23, 1984, pp. 832-835.

[2] Åqvist, L., and Ph. Mullock, *Causing Harm – A logical-legal study* , Berlin, New York, 1989.

[3] Borst, P., and H. Akkermans, 'Engineering Ontologies' in: *International Journal of Human Computer Interaction* ,Vol. 46, 1997, pp. 365-408.

[4] Breuker, J.A., A. Valente and R. Winkels, 'Legal Ontologies: a Functional View', in: P.R.S. Visser and Winkels, Editors, *Legal Ontologies*, pages 23-36, AC, New York, 1997.

[5] Hart, H.L.A., and T. Honoré, *Causation in the Law* , Oxford, Oxford University Press, 1985.

[6] Hulswit, M., *A Semeiotic Account of Causation – The 'Cement of the Universe' from a Peircean Perspective*, Nijmegen, form of Ph.D. thesis, 1998.

[7] Mackey, J.L., *The Cement of the Universe – A Study of Causation*, New York, Oxford University Press, 1986.

[8] McCarthy, J., and P. Hayes, 'Some Philosophical Problems from the Standpoint of Artificial Intelligence', in: *Machine Intelligence*, Vol. 4, 1969.

[9] Pearl, J., *Causality*, Cambridge (UK), Cambridge University Press, 2000

[10] Shoham, Y., *Reasoning about Change*, Cambridge, MIT Press, 1988.

[11] Valente, A., *Legal Knowledge Engineering – A modeling approach* , Amsterdam, IOS Press, 1995

[12] Winkels, R., A.W.F. Boer, J.A. Breuker and D. Bosscher, 'Assessment Based Legal Information Serving and Cooperative Dialogue in CLIME' , in: J.C. Hage, T.J.M. Bench-Capon, A.W. Koers, C.N.J. de Vey Mestdagh, C.A.F.M. Grütters (eds.), *Proceedings of Legal Knowledge Based Systems – JURIX: The Eleventh Conference*, Nijmegen, Gerard Noodt Instituut, 1998.