CARTWRIGHT'S THEOREM AND PROCEDURAL APPROACH TO CAUSALITY

Paweł Kawalec

Abstract

N. Cartwright's recent results on invariance under intervention and causality (2003) are reconsidered. Procedural approach to causality elicited in this paper and contrasted with Cartwright's apparently philosophical one unravels certain ramifications of her results. The procedural approach seems to license only a constrained notion of intervention and in consequence the "correctness to invariance" part of Cartwright's first theorem fails for a class of cases. The converse "invariance to correctness" part of the theorem relies heavily on modeling assumptions which prove to be difficult to validate in practice and are often buttressed by independently acquired evidence.

1. Introduction.

Recently N. Cartwright (2003) published results aiming to pin down the conceptual bonds

between invariance under intervention and causal dependence. The results, which are

commonly taken for granted, pertain to theories of causality. Besides, as Cartwright is right to

emphasize "[t]his project is important to practicing science" (2003, 205). This seems to

suggest that her primary focus is on the issues of causal inference as contrasted with the nature

of causality. But from the point of view of the former, as will be argued in the ensuing, her

results appear to have certain ramifications.

      The primary practical concern as regards causality, I take it, seems to be whether the

propounded causal claim can be sustained given the evidence. Fortunately, we don't have to

sit and wait until the relevant piece of evidence hits upon us, we reach for it by intervening on

a causal mechanism operating out there in the world, or so we believe. Drawing upon a causal

model at hand we set ourselves to form expectations as to what would happen given an action

is performed and then update the model to accord with the actual result of the action

performed. This rough picture seems to pick out how we intuitively relate interventions to

causality (cf. Woodward 2002, S369-S371) and in Section 3 below I will try to cast those

intuitions in a more precise form.

      And this is how I proceed in the ensuing. First, I highlight two common motivations

for drawing the link between causality and intervention. Next, I bring out the points of

difference between how the two characterize intervention. With the modified notion at hand I

will reconsider Cartwright's first theorem. The paper concludes with discussion of the plausibility of the assumptions of the theorem as confronted with "practicing science".

2. Intervention and Causality.

I proceed now to indicate distinctive features of the twofold relationship between causality and intervention. The pragmatic concern with intervention as underlying causality derives, at least in part, from the fact that in some cases it seems to be indispensable in discriminating which correlations among the data manifest a genuine causal relationship.[1] As D. Hausman and J. Woodward put it:

> [f]or most scientists, the crucial difference between the claim that X and Y are correlated and the claim that X causes Y is that the causal claim, unlike the claim about correlation, tells one what would happen if one intervened and changed the value of X. It is this feature of causal knowledge that is so important to action. (1999, 533)

This point on the relationship between causality and intervention may further be elaborated to back up scientists in inferring causal facts from data; such an approach will be henceforth referred to as *procedural*.[2] Alternatively, the elaboration may also be oriented towards solving conceptual issues pertaining to *philosophical* theories of causality, especially the issue of reductive analysis of this notion. The difference between procedural and philosophical approaches in discussing the relation between causality and intervention is succinctly and perspicuously worded by P. Menzies and H. Price thus

we disavow any view which states that the meaning of causal statements is to be given in terms of their verification-conditions or the conditions for their warranted assertion. In our view, the notion of agency enters the analysis of causation not because experimental manipulations provide evidence for causal claims, but because the notion of causation is an extrinsic one which by its very nature is rooted in the idea of manipulation. (1993, 196)

Hausman and Woodward emphasize that the discrepancy between the two approaches becomes apparent when one turns to a reductive analysis of the notion of causality:

A number of philosophers, including R. G. Collingwood, Douglas Gasking, G. H. von Wright, and most recently, Peter Menzies and Huw Price have defended a more ambitious thesis. They have held, not just that there is a connection between causation and manipulation … but that one can appeal to this connection to provide a reductive analysis of causation. (1999, 534; bibliographical references omitted)

Hausman and Woodward contrast then what I call here procedural with philosophical approach to causality, especially with agency theories, arguing that intervention "could not possibly provide a reductive analysis of causation" (1999, 535; Hausman 1997, S19-S23). Admittedly, I do not need to counter this claim while still holding that what is characteristic of philosophical theories of causality is that they take the notion of intervention to be constitutive of a notion of causality independently of how the meaning of the latter notion or a proposition in which it occurs is verified.

Of course, the above characterization is far from complete. The highlighted distinctive features of the two approaches will suffice, however, for the task at hand, i.e. to set apart two different notions of intervention: Cartwright's and procedural (as defined in Section 3 below).

Let me now turn to Cartwright's analysis of intervention, which – despite the appearance to the contrary – bears the mark of philosophical approach. In contrast to agency theories, Cartwright rejects outright to attribute a substantive meaning to intervention as "setting the target feature were we *wish* it to be, or *at will*, or *arbitrarily*" (2003, 204). For, as she argues, "[n]either set of theorems requires a notion so strong. All that is required is that nature allow specific kinds of *variation* in the features under study" (2003, 204), "[or] – she continues in the footnote – if the right kind of variation does not actually occur, there must be a fact of the matter about what would happen were it to do so" (2003, 204 f. 1).[3]

She acknowledges that intervention "of the right sort" may well suffice for the right kind of variability, while not the reverse. For intervention may "mislead us into demanding stronger tests for causality than are needed" (2003, 205).

The point is reaffirmed in Section 2 of her paper in the context of the notion of intervention required for the first theorem. She admits:

> it is the occurrence of these [causal] systems that matters for testing the correctness of causal claims; it is not necessary that they come to occur through anything naturally labeled an *intervention* or a *manipulation*. I shall, therefore, talk not of *intervention* but rather, of *variation*. (2003, 207)

And again in the footnote she continues:

> it is mistaken to argue that manipulation is central to causation *on the grounds that* one important kind of test for causal correctness – the "invariance" test – cannot do without it. (2003, 207 f. 5)

> Thus, Cartwright leaves room for discussion of the conceptual bond between manipulation and causation and therefore implicitly validates the distinction between philosophical and procedural approaches.

> The claims brought out in the above quotations coupled with the commitment that [t]his project is important to practicing science. When we know necessary or sufficient conditions for a causal-law claim to be correct, we can put them to use to devise real tests for scientific hypotheses. (2003, 205)

give the impression that hers bears the mark of a procedural notion of intervention.

On the definition she propounds (for a quote see Section 4) intervention takes place in a causal *system*, which is contrasted with its model, i.e. a causal *structure*. In the case a causal inference is pursued, as Woodward (2000, 204-205 f. 1) points out, causal systems fail to provide an "independent alternative (and non-trivial) characterization" of causal dependencies. Only when the issues pertaining to inferring causal facts from data are disregarded for the sake of conceptual bonds between the notion of intervention and causality that causal systems are put back to work. Therefore, the notion of intervention Cartwright propounds is designed to resolve conceptual issues pertaining to the nature of causality.[4]

3. Marking out Intervention.

Admittedly, a more detailed procedural theory relating intervention to causality will turn on

how the former notion is characterized. It has to be acknowledged that Cartwright's

remarkable contribution is in demonstrating that intervention comes in a variety.[5]

Nevertheless, a comprehensive outline of the whole territory seems not yet forthcoming. Even

her own definitions of intervention are, as she admits, fragmentary. So, I am not here going to

set out myself with this formidable task and will only focus on what sets the procedural notion

of intervention apart from others.

Cartwright's criticism (2002) of the characterization of intervention by Hausman and

Woodward (1999) elicits two kinds of intervention of major importance. With regard to a

causal system $\mu = <V, L>$, where V is a set of m variables and L is a set of causal laws

governing them, intervention may be set out either as modifying the laws themselves leading

thus to a modified causal system, or as changing the values of variables within the causal

system and examining the effects it has for other variables. In this paper I am concerned with

the former kind of intervention.

In response to Cartwright's criticism Hausman and Woodward present an explicit

definition of "an arrow-breaking" intervention (2004, 149-50). Roughly, in the deterministic

case an arrow-breaking intervention causes a given variable X_i to take on a particular value

"regardless of the values" of any other variables in this model as well as error terms and

interventions on other variables. The arrow-breaking intervention on X_i is idealized not to be causally involved with any other interventions or error terms and cause variables other than X_i only via this variable (Spirtes et al. 1999, 74). In other words, the arrow-breaking intervention is "a completely exogenous direct cause of X_i which, for some values, sets the value or probability distribution of X_i without regard to the values of the other causes of X_i" (2004, 150).

Both Cartwright's notion of intervention as well as Hausman and Woodward's raise concern. Let us turn to the former. As pointed out in Section 2, on this account intervention takes place in the causal system, which is contrasted with its model, i.e. the causal structure. Correctness of a given "causal-law" claim in causal structure is tested against the results of intervention as specified in the corresponding causal system. In the case of causal inference there is no independent characterization of causal system other than investigating causal structures. If causal correctness of a given claim is tested against pre-intervention and post-intervention distributions, then causal structure may not be uniquely identified. Therefore the procedural definition of intervention needs to face this problem.

Now, the definition of arrow-breaking intervention on a variable X_i formulated by Hausman and Woodward is explicitly related to the set of variables $\mathbf{V}$ to which X_i belongs. Moreover, it is also implicitly related to at least another part of a causal model, i.e. to the causal laws binding the variables in $\mathbf{V}$. This becomes apparent once we insist on a more explicit formulation of "regardless" which occurs in the definition (as quoted above). In

consequence at least two different types of models have to be accounted for – in Cartwright's terms – intervention takes place in a "test system" which is a transformation of a given "causal system".[6] Given the fact that in most cases neither of these is uniquely determined by the underlying distributions, the procedural definition of intervention refers to sets of observationally indistinguishable models[7] which are compatible with the distributions. [8]

To conclude the above remarks, I proceed to characterize the procedural notion of intervention. A *causal structure* $D = <\mathbf{V}, \mathbf{L}>$ consists of the set of variables $\mathbf{V}$ and the set of their functional dependencies $\mathbf{L}$ of the form $X_i = f_i(\mathbf{PA}_i, U_i)$ – where $\mathbf{PA}_i \subset \mathbf{V}$ is a set of parents of X_i (Pearl 2000, 14) and U_i is a random disturbance. A *latent structure* LS is a pair $<D, \mathbf{O}>$ (where $\mathbf{O} \subseteq \mathbf{V}$ is a set of observed variables). Now, one latent structure LS $= <D, \mathbf{O}>$ is *preferred* to another LS' $= <D', \mathbf{O}>$ iff for every set of parameters Θ_D compatible with D there exists $\Theta'_{D'}$ such that $P_{[O]}(<D', \Theta'_{D'}>) = P_{[O]}(<D, \Theta_D>)$. Finally, two further notions needed are: minimal LS and LS consistent with a distribution. LS is *minimal* with regard to a class of latent structures iff no member of that class is preferred to LS. And LS $= <D, \mathbf{O}>$ is *consistent* with a distribution P* over $\mathbf{O}$ if there exists a parametrization Θ_D such that $P_{[O]}(<D, \mathbf{O}>) = $ P*.

In general, the minimal set of structures consistent with a distribution P will extend beyond the true causal model. For a pre-intervention, i.e. nonexperimental distribution may leave unoriented links between some variables, may lead to wrong estimates of coefficients and may also leave out common causes (to be replaced by additional variables).

(PI) $\text{PI}_{\mathbf{X}}$: $\mathcal{L} \to \mathcal{L}_{\mathbf{X}}$ is an intervention on $\mathbf{X} \subseteq \mathbf{V}$ in LS $\in \mathcal{L}$ iff $\mathcal{L}$ is a set of minimal latent

structures consistent with (pre-intervention) stable distribution P over $\mathbf{V}$ and $\mathcal{L}_{\mathbf{X}}$ is a set

of minimal latent structures consistent with each (post-intervention) stable distribution

$P_{X_i} \in \mathbf{P_X}$ (where P_{X_i} is a transformation of P resulting from setting $X_i = x_i$ and dropping

all laws of the form $X_i = f_i(\mathbf{PA}_i, U_i)$; $\mathbf{P_X}$ is the set of such transformations for all

X_i in $\mathbf{X}$).

One idealization is that the distributions are stable (Pearl 2000, 48). Let I(P) denote the

set of all conditional independence relationships embodied in P. $I(P(<D, \Theta_D>))$ is *stable* iff

$I(P(<D, \Theta_D>)) \subseteq I(P(<D, \Theta'_D>))$.

Another idealization is that any constraints on which variables in $\mathbf{V}$ could be

intervened on are ignored and it is taken for granted that for any variable in the model one

could come up with an appropriate intervention. The upshot of these idealizations as with

respect to the "invariance to correctness" part of the proof of the first Cartwright's theorem

will be considered in Section 6.

In their joint paper Hausman and Woodward (1999) observe a subtle point as regards

observationally indistinguishable causal models. Admittedly, any intervention in a causal

system is modal. Even if two causal systems are observationally indistinguishable,

intervention may well sustain empirically verifiable counterfactuals yielding different

expectations. Apparently, a definition of intervention which aims at restoring our intuitive

understanding of this notion should not be limited to the specification of the actions involved,

but also pay close attention to its modal aspect. Such a definition should then tell us more about counterfactual states of affairs. Apparently, the definition (PI) satisfies this postulate as causal structures intervened on can vary without violating observational equivalence.

In order to align it with the procedural definition of intervention (PI) invariance needs to be recast as follows:

(I)   A functional dependence G: $X_i = f_i(\mathbf{PA}_i, U_i)$ in LS $\in \mathcal{L}_\mathbf{X}$ is invariant under intervention

$PI_\mathbf{PA}$: $\mathcal{L} \rightarrow \mathcal{L}_\mathbf{PA}$ $(\mathbf{PA} \subseteq \mathbf{PA}_i)$ iff G $\in \mathbf{L}'$ for every $\mathbf{L}'$ in $\mathcal{L}_\mathbf{PA}$.

One idealization endorsed in (I) is that the new distributions in $\mathbf{P}_\mathbf{PA}$ change only those dependencies that could not be read off pre-intervention distribution P.

4. **Cartwright's Theorems on Causality and Invariance.**

Cartwright characterizes the task she sets herself out to accomplish succinctly thus:

> To get clear about whether invariance under intervention is or is not necessary or sufficient for a causal-law claim to be correct, and under what conditions, we need to know what counts as an intervention, what invariance is, and what it is for a causal-law claim to be correct. (2003, 204)

Both sets of theorems apply to linear deterministic systems. The first one may roughly be characterized as concerned with modification of a causal mechanism which leads to adding and dropping the relevant causal laws within a given causal system.

So, Cartwright's first theorem reads:

A functionally true equation is causally correct iff it is invariant in all its independent

variables, either singly or in any combination. (2003, 208)

Intervention on q in a given causal system µ results in a "test system", which appends the set

of laws in µ with the causal law that a target quantity q is set to a specific value Q (within the

allowed range of q) and, as a consequence, all laws which take q as an effect are dropped.[9]

Cartwright defines the relevant kind of invariance thus:

An equation in a (linear deterministic) causal system <Q, CL> giving a true functional

relation (but not necessarily one that replicates one of nature's causal laws) is *invariant*

in q iff it continues to give a true functional relation for any value that q takes in any

test situation[10] for q relative to <Q, CL>. (2003, 208)[11]

Admittedly, the proof of the implication "correctness to invariance" turns on the notion

of invariance – on Cartwright's definition this holds trivially for the only kind of laws dropped

in the test system are those where q_i was taken as a function of its causal parents. Therefore, as

one would expect, this result may not hold for a different notion of intervention.

5. Variance and Effective Intervention.

The procedural approach endorses tailoring of both the notion of intervention as well as

invariance. The definition of intervention (PI) and of invariance (I) seem to incorporate the

proper adjustments. With the modified notions I proceed to re-examine separately both

implications of Cartwright's first theorem.

Correctness to Invariance

For all H: $X_i = f_i(\mathbf{PA}_i, U_i)$ if H is in a LS (compatible with distribution P), then H is invariant under intervention $\mathrm{PI}_{\mathbf{PA}_i}$ in all $X_j \in \mathbf{PA}_i$.

Consider LS in which G_2: $X_i = f_i(X_j, U_i)$ holds. Suppose that in the set of variables $\mathbf{V}$ in LS no X_k exists such that it would yield G_3: $X_k = f_k(X_i, X_j, \ldots, U_i)$. It is then possible to find a distribution P_{X_j} such that X_j is dependent on X_i, but not the reverse.

Apart from this particular class of cases this implication holds analogously to the original.

Invariance to Correctness

For all H: $X_i = f_i(\mathbf{PA}_i, U_i)$ if H is invariant under intervention $\mathrm{PI}_{\mathbf{PA}_i}$ in all $X_j \in \mathbf{PA}_i$, then H is in a LS (compatible with distribution P).

Proof. The intervention $\mathrm{PI}_{\mathbf{PA}_i}$ transforms P into the space of probability distributions $\mathbf{P}_{\mathbf{PA}_i}$. It follows from the definition of $\mathrm{PI}_{\mathbf{PA}_i}$ that for all $\mathbf{L'} \in \mathcal{L}_{\mathbf{PA}_i}$ $H \in \mathbf{L'}$. H preserves then all probabilistic independencies for all $P_{X_n} \in \mathbf{P}_{\mathbf{PA}_i}$. Hence, H is in at least one observationally indistinguishable LS compatible with P.

In the next section we need to face reasons why the assumptions and idealizations used in establishing the latter implication could – and sometimes are – undermined in "practicing science".

6. Causal Systems and Confounders.

We may now ponder to what extent the inference from "invariance to correctness" is a matter of conceptual analysis and if the stubbornness in applying it is not an argument for its contingency.

Pearl, for one, answers succinctly thus:

the assumption that one is in the possession of a probability function on all variables relevant to a given domain … absolves the analyst from worrying about unmeasured spurious causes that might (physically) affect several variables in the analysis and still remain obscure to the analyst. It is well known that the presence of such "confounders" may reverse or negate any causal conclusion that might be drawn from probabilities. (2000, 252)

The most ardent critic of straightforward acceptance of "invariance to correcntess" on merely conceptual grounds, David Freedman, gives a list of reasons why the assumptions fail in many empirical studies:

Effects are weak and inconsistent. Endpoints are poorly defined. There is an extensive search for statistical significance. Important confounders are ignored. (1999, 12)[12]

A compelling example of successful scientific investigation on causal facts he cites is research on cervical cancer:

> This cancer has been studied for many years. Some investigators have identified the cause as tissue irritation; others point to syphilis, or herpes, or chlamydia; still others have found circumcision of the husband to be protective. … Today it is believed that cervical cancer is in large part a sexually transmitted disease, the agent being certain types of human papillomavirus, or HPV. There is suggestive evidence for this proposition from epidemiology and from clinical practice, as well as quite strong evidence from molecular biology. If so, the earlier investigators were mislead by confounding. (1999, 13)

> The moral is then that in establishing causal influence one needs to combine different kinds of evidence in order to get an independent warrant in case the modeling assumptions turn out defective. Commenting on the result that smoking causes lung cancer Freedman concedes:

> The strength of the case rests on the size and coherence of the effects, the design of the underlying epidemiologic studies, and on replication in many contexts. Great care was taken to exclude alternative explanations for the findings. Even so, the argument requires a complex interplay among many lines of evidence. Regression models are peripheral to the enterprise. (1999, 16)

Analogous observations are ubiquitous in the literature and often are stressed by scientists who deal with weak causal dependencies, especially sociologists and epidemiologists.[13]

These critical remarks, especially Freedman's objections (Freedman and Humphreys 1999), could be alleviated by paying attention to disparities in investigated cases. In a recent paper A. Gopnik, C. Glymour and their co-workers (2004) compellingly argue that apart from data a system for causal inference draws upon two kinds of assumptions: formal (general knowledge mainly concerning mathematical properties of data) and substantive (domain specific knowledge about types of causal dependencies).[14] It might then be conjectured that in cases when substantive knowledge is considerable, the modeling assumptions could be relied upon with a confidence (Pearl 1995b, 12). Otherwise, researchers have to gather more varied and more ample evidence independently of the modeling assumptions:

> [i]n complex fields like the social sciences and epidemiology, there are only few (if any) real life situations where we can make enough compelling assumptions that would lead to identification of causal effects. (Pearl personal communication, I quote from Freedman 2002, 15)

7. Conclusion.

A proponent of philosophical approach need not be distracted by this kind of problems for in this case the "invariance to correctness" part of Cartwright's theorem may rely on

assistance of substantial metaphysical claims. A prominent illustration is the notion of invariance understood as stable configurations of capacities which give rise to "nomological machines" manifesting themselves in causal laws.[15] Only when confronted with the problem of identification – pertaining to procedural approach – this kind of solution will perhaps drive us back to the discussion of the previous section.

To conclude, it turns out that when we attend to the problems of "practicing science" we do not seem to be driven by Cartwright's results far beyond her dictum "No causes in, no causes out".

References

Cartwright, Nancy (1999), *The Dappled World: A Study of the Boundaries of Science.*
Cambridge: Cambridge University Press.

⸺ (2002), "Against Modularity, the Causal Markov Condition and Any Link Between the
Two: Comments on Hausman and Woodward", *British Journal for the Philosophy of
Science* 53: 411-453.

⸺ (2003), "Two Theorems on Invariance and Causality", *Philosophy of Science* 70: 203-
224.

Cooper, Gregory F. (1999), "On Overview of the Representation and Discovery of Causal
Relationships Using Bayesian Networks", Glymour and Cooper 1999, 3-62.

Freedman, David A. (1999), "From Association to Causation: Some Remarks on the History
of Statistics", Technical Report 521, Department of Statistics, Berkeley: University of
California.

⸺ (2002), *On Specifying Graphical Models for Causation, and the Identification Problem*.
Technical Report 601, Department of Statistics, Berkeley: University of California.

Freedman, David A., and Paul Humphreys (1999), "Are There Algorithms That Discover
Causal Structure?", *Synthése* 121: 29-54.

Glymour, Clark, and Gregory F. Cooper (eds.) (1999), *Computation, Causation, and
Discovery.* Menlo Park, CA: AAAI Press, Cambridge, MA: MIT Press.

Gopnik, Alison, Clark Glymour, David M. Sobel, Laura E. Schulz, Tamar Kushnir, and David

Danks (2004), "A Theory of Causal Learning in Children: Causal Maps and Bayes Nets",

Psychological Review 111: 3-32.

Hausman, Daniel M. (1997), "Causation, Agency, and Independence", *Philosophy of Science*

(Proceedings) 64: S15-S25.

—— and James Woodward (2004), "Modularity and the Causal Markov Condition: A

Restatement", *British Journal for the Philosophy of Science* 55: 147-161.

Menzies, Peter, and Huw Price (1993), "Causation as a Secondary Quality", *British Journal*

for the Philosophy of Science 44: 187-203.

Pearl, Judea (1995a), "Causal Diagrams for Empirical Research", *Biometrika* 82: 669-710.

—— (1995b), "Causal Inference from Indirect Experiments", *Artificial Intelligence in Medicine*

7: 561-582.

—— (1999), "Graphs, Structural Models, and Causality", in Glymour and Cooper 1999, 95-

138.

—— (2000), *Causality: Models, Reasoning, and Inference.* Cambridge: Cambridge University

Press.

Psillos, Stathis (2002), *Causation and Explanation*. Montreal: McGill-Queen's University

Press.

Robins, James M. (1999), "Testing and Estimation of Direct Effects be Reparametrizing Directed Acyclic Graphs with Structural Nested Models", Glymour and Cooper 1999, 349-405.

Spirtes, Peter, Clark Glymour, and Richard Scheines (2000), *Causation, Prediction, and Search*, 2nd. edition. Cambridge, MA: MIT Press.

Spirtes, Peter, Clark Glymour, Richard Scheines, Christopher Meek, Stephen Fienberg, and Elizabeth Slate (1999), "Prediction and Experimental Design with Graphical Causal Models", Glymour and Cooper 1999, 65-93.

Weber, Max (1922), *Gesammelte Aufsätze zur Wissenschaftslehre*. Tübingen: Mohr-Siebeck.

Woodward, James (1999), "Causal Interpretation in Systems of Equations", *Synthése* 121: 199-247.

⸺ (2000), "Explanation and Invariance in the Special Sciences", *British Journal for the Philosophy of Science* 51: 197-254.

⸺ (2002), "What is a Mechanism? A Counterfactual Account", *Philosophy of Science* (Proceedings) 69: S366-S377.

[1] Especially if it is neither the case when a causal dependence is uniquely determined by a distribution, nor the case when substantive assumptions could assist scientists in this task (Pearl 2000, 19; Gopnik et al., 2004).

[2] A perspicuous illustration of this approach is in Cooper 1999, 9-11.

[3] As emphasized by Woodward (1999, 202) what matters for this kind of intervention is not the result of actually performed experiment, but the result of an ideal experiment which perfectly conforms to assumptions underlying the causal system on which the calculation of the causal effect of intervention is performed.

[4] In commenting on Woodward's claim S. Psillos observes a distinction analogous to the one postulated here between procedural and philosophical approaches to causality, but questions that the latter may also apply to the notion of intervention (2002, 106).

[5] Hausman and Woodward (2004) seem to counter this claim for its purported entailment of causal pluralism. But that such an entailment prevails would require a separate argument. See also Woodward 1999, 202.

[6] M. Weber, for one, claimed that the assessment of causal powers of historical and unrepeatable events is only possible when a set of alternative models is considered rather than a particular model with the event considered (1922, 287).

[7] A similar point is discussed in Spirtes et al. 2000, 50. Moreover, Spirtes et al. introduce

variety of notions corresponding to "observational indistinguishability" (2000, 59-72).

[8] For a precise formulation see Pearl 2000, 24.

[9] This may come in an explicit form as well as be derived by numerical transitivity (A_4), see

Cartwright 2003, 207.

[10] I assume that the phrase "test situation" in the definition should read "test system".

[11] There is an ambiguity in describing the first element in <Q, CL> as an ordered set of

quantities $<q_1, \ldots, q_m>$. The second element is a set of causal laws q_k c= $\Sigma_{j<k}a_{kj}q_j$, where a_{kj} is a

real number and the laws satisfy Cartwright six axioms (2003, 205-206). Moreover, the notion

of invariance in the proof is ambiguous as Cartwright moves between invariance of causal

systems and – an undefined concept of – invariance of representation of those systems.

[12] For a similar discussion and list see also Cooper 1999, 11.

[13] J. Robins comes up with a telling argument. He identifies a case in which a hypothesis that a

given variable has a direct effect on another variable is identifiable based on the observed data,

but "cannot be tested simply by testing for the presence of absence of arrows in the DAG"

(1999, 350).

[14] A similar distinction is observed also in Pearl 1995a, 1999.

[15] Cartwright's focus on what "capacities do and why we need them" rather than on what they are obscures the fact that her interest is ultimately ontological (1999, 73; Psillos 2002, 192). This problem seems analogous to the one encountered in Section 2 in classifying Cartwright's approach to causality as procedural.