8 Cartwright's Realist Toil
From Entities to Capacities[1]

Stathis Psillos

INTRODUCTION

Nancy Cartwright has been both an empiricist and a realist. Where many philosophers have thought that these two positions are incompatible (or, at any rate, very strange bedfellows), right from her first book, the much-discussed and controversial *How the Laws of Physics Lie*, Cartwright tried to make a case for the following view: if empiricism allows a certain type of method in its methodological arsenal (inference to the most likely cause), then an empiricist cannot but be a scientific realist—in the metaphysically interesting sense of being ontically committed to the existence of unobservable entities. Many empiricists thought that because empiricism has been traditionally antimetaphysics, it has to be antirealist. One of the major contributions that Cartwright has made to philosophy of science is, I think, precisely this: there is a sense in which metaphysics can be respectable to empiricists. Hence, scientific realism cannot be dismissed on the grounds that it ventures into metaphysics. To be sure, the metaphysics that Cartwright is fond of is not of the standard a priori (or armchair) sort. It is tied to scientific practice and aims to recover basic elements of this practice (e.g., causal inference). But it is metaphysics, nonetheless.

Cartwright's realism has been described as "entity realism". This is not accidental. She has repeatedly made claims such as 'I believe in theoretical entities' (Cartwright 1983: 89, see also 92). Typically, she contrasts her commitment to entities to her denial of "theoretical laws". In the sections 'Causal explanation' and 'Causal inference', I examine in some detail the grounds on which Cartwright tried to draw a line between being committed to entities and being committed to theoretical laws, and I find them wanting. In 'Causal inference' I also claim that the method Cartwright articulated, Inference to the Most Likely Cause, is important but incomplete. Specifically, I claim that there is a more exciting method that Cartwright herself describes as Inference to the Best Cause, which, however, is an instance, or a species of Inference to the Best Explanation. But Cartwright has been against Inference to the Best Explanation (IBE). So, in the section 'Why deny inference to

the best explanation?' I consider and try to challenge Cartwright's central argument against IBE.

At least part of the motivation for her early, restricted, realism was a certain understanding of what scientific realism is. She took scientific realism to entail the view that the world has a certain hierarchical structure, where the more fundamental laws explain the less fundamental ones as well as the particular matters of fact. In *The Dappled World*, she rightly disentangled these issues. 'Nowadays', she says, 'I think I was deluded by the enemy: it is not *realism* but *fundamentalism* that we need to combat' (Cartwright 1999: 23). What, I think, emerges quite clearly from her later writings is that Cartwright does not object to realism. Rather, she objects to Humeanism about laws, causation, and explanation. Insofar as Humeanism is a metaphysics independent of scientific realism, Cartwright is a more full-blown realist, without being Humean. And this is what she is. In the penultimate section, 'Capacities', I discuss in some detail Cartwright's central non-Humean concept, viz., capacities. Cartwright is a strong realist about capacities. They are the fundamental building blocks of her metaphysics. But there seem to be a number of problems with capacities. Though we can easily see how attractive it is to be a realist about capacities, I think it's really hard to be one. So, though Humeanism is certainly independent of scientific realism, I argue that we have not been given compelling reasons for a non-Humean metaphysics of capacities.[2]

It is helpful to state clearly five worries about Cartwright's views that I develop in this paper. The first is that though she was right to insist on the ontic commitment that flows from causal explanation, she was wrong to tie this commitment solely to the entities that do the causal explaining. This move obscured the nature of causal explanation and its connection to laws. The second worry is that when she turned her attention to causal inference, by insisting on the motto of "the most likely cause", she underplayed her powerful argument for realism. For she focused her attention on an extrinsic feature of causal inference (or, indeed, of any ampliative inference), viz., the demand of high probability, leaving behind the intrinsic qualities that causal explanation should have in order to provide the required understanding. The third is that her objections to Inference to the Best Explanation were unnecessarily tied to her objections about the falsity of fundamental laws. Fourth is that though her argument for positing capacities and being realist about them was supposed to take strength from its parallel with Sellars's powerful argument for the indispensable explanatory role of positing unobservable entities, there are important disanalogies between the two arguments that cast doubt on the indispensability of capacities. The final (fifth) worry is that laws—perhaps brute regularities—might well have to come back from the front door, as they are still the most plausible candidates for explaining why objects have the capacities to do what they can do.

CAUSAL EXPLANATION

One of Cartwright's central claims is that causal explanation is ontically committing to the entities that do the explaining (Cartwright 1983). Here are some typical statements of it:

> That kind of explanation succeeds only if the process described actually occurs. To the extent that we find the causal explanation acceptable, we must believe in the causes described (Cartwright 1983: 5).
>
> In causal explanations truth is essential to explanatory success (Cartwright 1983: 10).
>
> But causal explanations have truth built into them (Cartwright 1983: 91).
>
> (. . .) existence is an internal characteristic of causal explanation (Cartwright 1983: 93).

These assertions are not all equivalent to one another, but I do not dwell on that. For, there is indeed something special with causal explanation. So, let's try to find out what it is. As a start, note that it is one thing to say that causal explanation is ontically committing but quite another thing to say what a causal explanation *is*. Let's take them in turn.

Ontic Commitment

If *c* caused *e*, then, clearly there must be events *c* and *e* which are thus causally connected. This follows almost directly from the standard Davidsonian account of singular causal statements. Causation is not quite the same as causal explanation, but causes do explain their effects, and there is, to say the least, no harm in saying that if *c* causes *e* then *c* causally explains *e*. This feature of causal explanation by virtue of which it is ontically committing to whatever does the causing is not peculiar to it. Compare the relation *c* preceded *e*: *c* must exist in order to precede *e*. So, Cartwright's claim is an instance of the point that the relata of an actual relation *R* must exist in order for them to be related to each other by *R*. I think this is what Cartwright should mean when she says that '(. . .) existence is an internal characteristic of causal explanation' (Cartwright 1983: 93).

An equivalent way to show that causal explanation is ontically committing is this. To say that the statement "*c* causally explains *e*" is ontically committing to *c* and *e* is to say that "*c* causally explains *e*" is true. This way of putting things might raise the spectre of van Fraassen, as Hitchcock reminds us (Hitchcock 1992). Couldn't one just accept that "*c* causally explains *e*"

without believing that it is true? And if so, couldn't one simply avoid the relevant ontic commitments to whatever entities are necessary to make this statement true? Indeed, insofar as we can make sense of an attitude towards a statement with a truth-condition which involves acceptance but not belief, van Fraassen is on safe ground here. He is not forced to believe in the truth of statements of the form "*c* causally explains *e*". Cartwright's point, however, is not meant to be epistemic. Her point is, I think, twofold. On the one hand, she stresses that we cannot avoid commitment to the things that are required to make our assertions true. On the other hand (and more importantly), insofar as we do make some assertions of the form "*c* causally explains *e*" (e.g., about observable events such as shortcircuits and fires or aspirins and headaches), there is no reason not to make others (e.g., about unobservable entities and their properties).

So, causal explanation is egalitarian: It sees through the observable–unobservable distinction. It is equally ontically committing to both types of entity, precisely because the relation of causal explanation is insensitive to the observability of its relata. In other words, what matters for ontic commitment is the causal bonding of the relata of a causal explanation. So, Cartwright's point is that there is just one way to be committed to entities (either observable or unobservable) and it is effected through causal explanation.

What Exactly Is a Causal Explanation?

This remains an unsettled question, even after it is accepted that causal explanation is ontically committing. The question, in a different form, is this: What exactly is the relation between *c* and *e* if *c* causally explains *e*? In the literature, there have been a number of attempts to explain this relation. I do not discuss them here.[3] Cartwright has offered a gloss of the relation *c causally explains e*. She put forward an early version of the contextual unanimity principle, viz., the idea that *c* causes *e* iff *c* increases the probability of *e* in all situations (contexts) which are causally homogeneous with respect to the effect *e* (Cartwright 1983: 25–26). I do not dwell on this principle here. But one thing is relevant. Although principles such as the above do cast some light on the notion of causal explanation, they do not offer an analysis of it, as they presuppose some notion of causal law or some notion of causally homogeneous situation. Cartwright is very clear on this when she says, for instance, that what makes the decay of uranium 'count as a good explanation for the clicks in the Geiger counter' is not the probabilistic relations that obtain between the two events 'but rather the causal law—"Uranium causes radioactivity"' (Cartwright 1983: 27). Still, it might be said that though Cartwright does not offer 'a model of causal explanation' (Cartwright 1983: 29), she does constrain this notion by objecting to certain features that causal explanation is taken to have. Most centrally, she objects to the deductive-nomological model of causal explanation. But it is not clear, for instance, that she takes a singularist account

of causal explanation. In fact, it seems that she doesn't. For she allows that certain 'detailed causal principles and concrete phenomenological laws' are involved in causal explanation (Cartwright 1983: 8). Her objection is about laws captured by 'the abstract equations of a fundamental theory' (Cartwright 1983: 8). So, even if she objects to the thesis that all causal explanation should be nomological, she doesn't seem to object to the weaker thesis that at least some causal explanation should be nomological. In any case, it's one thing to deny that the laws involved in causal explanation are the abstract high-level laws of a theory and it is quite another to deny that laws, albeit low-level ones, are involved in, or ground, causal explanation. For all I know, Cartwright does not deny the latter (Cartwright 1983).

Here is the rub, then. If laws are presupposed for causal explanation, then it's no longer obvious that in offering causal explanations we are committed just to the relata of the causal explanation. To say the least, we should also be committed to a Davidson-style compromise that there are laws that govern the causal linkage between cause and effect. Though these laws might not be stateable or known, they cannot be eliminated. But this is not the end of it. Considering Davidson's idea, Hempel noted that when the existence of the law is asserted but the law is not explicitly stated, the causal explanation is comparable to having 'a note saying that there is a treasure hidden somewhere' (Hempel 1965: 349). Such a note would be worthless unless 'the location of the treasure is more narrowly circumscribed'. Think of it as advice: where there is causal explanation, search for the law that makes it possible. It's a side issue whether this law is a fundamental one or a phenomenological one or what have you. This is a worry about the kinds of law there are and not about the role of laws in causal explanation.

So here is my first conclusion. Cartwright's advertised entity-realism underplays her important argument for ontic commitment. In offering causal explanations, we are committed to much more than entities. We are also committed to laws, unless of course there is a cogent and general story to be told about causal explanation that does *not* involve laws. Note that it is not a reply to my charge that there might be a singular causal explanation. This is accepted by almost everybody—given the right gloss on what it consists in. Nor would it be a reply to my charge that, occasionally, we do not rely on laws to offer a causal explanation. A suitable reply would have to show that causal explanation is totally disconnected from laws. This kind of reply might be seen as being offered by Cartwright when she introduces capacities. But, as we shall see in the section 'Capacities', it is at least questionable that we can make sense of capacities without reference to laws.

CAUSAL INFERENCE

Given the centrality of causal explanation in Cartwright's argument for realism, one would have expected her to stay firmly in the business of explaining

its nature. But Cartwright does something *prima facie* puzzling. She spends most of *How the Laws of Physics Lie* (1983) on an attempt to cast light on the nature of the inference that takes place when a causal explanation is offered and on the conditions under which this inference is legitimate. (Doesn't that remind you of what Hume did?) One way to read what Cartwright does is this: she is concerned with showing when a potential causal explanation can be accepted as the actual one. More specifically, she is concerned with showing that there is something special in causal explanatory inference that makes it sound (or, at any rate, makes it easier to check whether it is sound or not). She says:

> Causal reasoning provides good grounds for our beliefs in theoretical entities. Given our general knowledge about what kinds of conditions and happenings are possible in the circumstances, we reason backwards from the detailed structure of the effects to exactly what characteristics the causes must have to bring them about. (Cartwright 1983: 6)

Thus put, causal reasoning is just a species of ampliative reasoning. From an epistemic point of view, that the explanation offered in this reasoning is causal (that is, that it talks about the putative causes of the effects) is of no special importance. What matters is what reason we have to accept the conclusion about the putative cause.

This seems to me a crucial observation. Cartwright explicitly draws a contrast between "theoretical explanation" and "causal explanation" (Cartwright 1983: 12). But this is, at least partly, unfortunate. For it obscures the basic issue at stake. *Qua* inferential procedures, causal explanation and theoretical explanation are on a par. They are each species of ampliative reasoning, and the very same justificatory problems apply to both of them (perhaps to a different degree).

Cartwright does think that there is something special in the claim that the inference she has in mind relies on a *causal* explanation. She calls this inferential process 'inference to the most likely cause' (Cartwright 1983: 6)—henceforth, IMLC. But there is a sense in which the weight is on the "most likely" and not on the "cause". It's just that Cartwright thinks that it's most likely to get things right if you are looking for causes than if you are looking for something else (e.g., general theoretical explanations). Before we see whether this is really so, let us press the present point a bit more.

Inference to the Most Likely Cause

What kind of inference is IMLC? An obvious thought is that we infer the conclusion (viz., that the cause is *c*) if and only if the probability of this conclusion is high. But this is a general constraint on any kind of ampliative inference with rules of detachment, and hence there is nothing special in IMLC in this respect. A further thought then might be that in the case

of IMLC there is a rather safe way to get the required high probability. The safety comes from relevant background knowledge of all sorts: that the effect has a cause, because in general effects do have causes; that we are offered a rather detailed story as to what the causal mechanism is and how it operates to bring about the effect; that we have controlled for all(?) other potential causes, etc. (Cartwright 1983: 6). All this is very instructive. However, thus described, IMLC gets its authority not as a special mode of inference where the weight is carried by the claim that *c* causally explains *e* but from whatever considerations help increase our confidence that the chosen hypothesis (viz., that it was *c* that caused *e*) is likely to be true. If these considerations are found wanting (if, for instance, our relevant background knowledge is not secure enough, or if we do not eliminate all potential alternative causes, or if the situation is very complex), then the claim that *c* causally explains *e* is inferentially insecure. It simply cannot be drawn, because it is not licensed as likely.

Indeed, my present complaint can be strengthened. Consider what Cartwright says: '(. . .) causal accounts have an independent test of their truth: we can perform controlled experiments to find out if our causal stories are right or wrong' (Cartwright 1983: 82). If we take this seriously, then all the excitement of IMLC is either lost or becomes parasitic on the excitement of a controlled experiment. It is lost if for every instance of an IMLC it is required that a controlled experiment is performed to check the conclusion of the inference independently. So, what if the excitement of IMLC becomes parasitic on the excitement of a controlled experiment? Controlled experiments are indeed exciting. But their excitement comes mostly from the fact that they are designed to draw safe causal conclusions, irrespective of whether there is on offer a causal explanation of the effect. When it is established by a clinical trial that drug *D* causes relief from symptom *S*, we may still be in the dark as to *how* and *why* this is effected, what the mechanisms are, what the detailed causal story is, etc. I think that causal explanation—*qua* inference—is exciting not just because we can get conclusions that are likely to be correct, but also because we get an understanding of *how* and *why* the effect is produced. But so far, we have got only (or mostly) the former. The hard question, I think, remains unaddressed: What is this (if anything) in virtue of which a causal explanation—*qua* an explanatory story—licenses the conclusion that it is likely to be correct? Put in more general terms, the hard problem is to find an intrinsic feature of causal explanation in virtue of which it has a claim to correctness and not just an extrinsic feature, viz., that there are independent reasons to think it is likely.

Inference to the Best Cause

Cartwright seems aware of the need for such an intrinsic feature. Occasionally, she describes IMLC as 'inference to the best cause' (Cartwright 1983: 85). I think this is not just a slip. Reference to "best cause" is not just meant

to *contrast* IBC to Inference to the Best Explanation (IBE), by replacing "explanation" with "cause". It is also meant, rightly I think, to *connect* IBC to IBE. It is meant to base the inference (the detachment of the conclusion) on certain features of the connection between the premises and the conclusion, viz., that there is a genuinely explanatory relation between the explanation offered and the explanandum. The "best cause" is not *just* a likely cause; it is a putative cause that causally explains the effect in the sense that it offers genuine understanding of how and why the effect was brought about. Cartwright says of Perrin's "best cause": 'we are entitled to [infer the existence of atoms] because we assume that causes make effects occur in just the way they do, via specific, concrete causal process' (Cartwright 1983: 85). If all we were interested in was high probability, then we wouldn't go for specific, concrete causal processes—for the more detail we put in, the more unlikely they become. The specific, concrete causal processes matter for understanding, not for probability.

The upshot is that if we conceive causal inference as Inference to the Best Cause (IBC), then it is no longer obvious that it is radically different from what has come to be known as Inference to the Best Explanation (IBE). The leading idea behind IBE—no matter how it is formulated in detail—is that explanatory considerations are a guide to inference. The inference we are concerned with is ampliative—and hence deductively invalid. But this is no real charge. Inferential legitimacy is not solely the privilege of deductive inference. IBC can then be seen as a species of IBE. It's a species of a genus, whose *differentia* is that in IBC the explanations are causal (see Psillos 2002b for details).

What sort of inference is IBE? There are two broad answers to this. (1) We infer to the probable truth of the likeliest potential explanation insofar as and because it is the likeliest explanation. On this answer, what matters is how likely the explanatory hypothesis is. (2) The best explanation, *qua* explanation, is likely to be true (or, at least more likely to be true than worse explanations). That is, the fact that a hypothesis H is the *best* explanation of the evidence issues a warrant that H is likely. The late Peter Lipton noted that the first answer views IBE as an inference to the Likeliest Potential Explanation, whereas the second views it as an inference to the Loveliest Potential Explanation (Lipton 1991: 61–65). The loveliest potential explanation is 'the one which would, if correct, be the most explanatory or provide the most understanding' (Lipton 1991: 61).

Exactly the same distinction applies to causal inference. If we think of it as an Inference to the Most Likely Cause (IMLC), then, as we have seen, the inferential weight is carried by the likeliness of the proposed causal explanation. So, it's not that a causal explanation is offered that licenses the inference. Rather, it is that this proposed explanation has been rendered likely. This rendering is extrinsic to the explanatory quality of the proposed explanation and relates to what we have done to exclude other potential explanations as likely. On the other hand, if we think of causal inference as Inference

to the Best Cause, we are committed to the view that the inferential weight is carried by the explanatory quality of the causal explanation offered, on its own *and* in relation to competing alternatives. Roughly put, the weight is carried by the understanding offered by the causal story and by the explanatory qualities that this story possesses.

Indeed, Cartwright speaks freely of "causal accounts" or "causal stories" offered by causal explanations. The issue then is not just to accept that there must be entities that make these causal accounts true. It is also to assess these accounts *qua* explanatory stories. If we take IBC seriously, there must be ways to assess these accounts, and these ways must be guides to whether we should accept them as true. It seems then that we need to take account of explanatory virtues (a) if we want to make IBC have a claim to truth; and (b) if we want to tie this claim to truth not just to extrinsic features of causal explanation (e.g., that it is more likely than other potential explanations) but also to intrinsic features of the specific causal explanatory story.

So, let me draw the conclusion of this section. Thinking of causal explanation as an inference to the best cause will require assessing the causal story offered, and this is bound to be based on explanatory considerations which align IBC to IBE.[4]

WHY DENY INFERENCE TO THE BEST EXPLANATION?

It is well known, however, that Cartwright resists IBE (Cartwright 1983). And it is equally well known that she thinks she is not committed to IBE, when she vouches for IBC. So the issue is by no means over. Cartwright explicitly denies that 'explanation is a guide to truth' (Cartwright 1983: 4) and discusses this issue quite extensively (Cartwright 1983). Due to lack of space, I focus on one of her arguments, which seems to me to be the most central one. This is the argument from the falsity of laws. But before I go into this, allow me to note an interesting tension in her current views on the matter.

The Transcendental Argument

Cartwright has always tried to resist global applications of IBE. In particular, she tried to resist versions of the "no miracles argument" for realism.[5] Consider her claim:

> I think we should instead focus on the causal roles which the theory gives to these strange objects: exactly how are they supposed to bring about the effects which are attributed to them, and exactly how good is our evidence that they do so? The general success of the theory at producing accurate predictions, or at unifying what before had been disparate, is of no help here. (Cartwright 1983: 8)

The last sentence of this quotation is, to say the least, overstated. But let's not worry about this now. For, in her current views, the general antitheory tone (Cartwright 1983) has been superseded by a more considered judgement about theories and truth. She concedes that 'the impressive empirical successes of our best physics theories may argue for the truth of these theories', but, as we have already seen, she denies that it argues 'for their universality' (Cartwright 1999: 4). In fact, her talk about 'different kinds of knowledge in a variety of different domains across a range of highly differentiated situations' implies that truth is in the vicinity. For knowledge without truth is an oxymoron. So, her objections to Inference to the Best Explanation do not try to challenge the very possibility of a link between explanation and truth. Rather, they aim to block gross and global applications of IBE.

Let us look at Cartwright's argument for "local realism", which, as she says, is supposed to be a Kantian transcendental argument (Cartwright 1999: 23). The way she sets it up is this: We have *X*—'the possibility of planning, prediction, manipulation, control and policy setting'. But without Φ—'the objectivity of local knowledge'—*X* would be impossible or inconceivable. Hence Φ. It's fully understandable why Cartwright attempts to offer a transcendental argument. These arguments are dressed up as deductive. Hence, they are taken not to have a problematic logical form. They compare favourably with IBE. But apart from general worries about the nature and power of transcendental arguments[6], there is a more specific worry: Is the above argument really deductive?

A cursory look at it suggests that it is: "Φ is necessary for *X*; *X*; Therefore, Φ". But it is misleading to cast it as above, simply because it is misleading to say that Cartwright's Φ is necessary for *X*. Kant thought that Euclidean geometry was necessary for experience. Of course, it isn't. He could instead have argued that *some* form of spatial framework is necessary for experience. This might well be true. But now it no longer deductively follows that Euclidean geometry must be true. In a similar fashion, all that Cartwright's argument could show is that something—call it Φ—is necessary for 'the possibility of planning, prediction, manipulation, control and policy setting'. But now, it no longer follows deductively that this Φ *must* be the realist's "objective local knowledge", no matter how locally or thinly we interpret this. To say the least, this Φ could be just empirically adequate beliefs, or unrefuted beliefs, or beliefs that the world cooperates only when we actually try to set plans, make observations, manipulate causes, etc. Put in a different way, all that follows from Cartwright's transcendental argument is a disjunction: Either objective local knowledge, or empirically adequate beliefs, or . . . is necessary for the possibility of planning, prediction, manipulation, control, and policy setting. But which disjunct is the true one? Further argument is surely necessary. There cannot be a transcendental deduction of objective local knowledge.

My suggestion is that the move from the "the possibility of planning, prediction, manipulation, control and policy setting" to a realist understanding of what needs to be the case for all of them to be possible (or, why not, actual) can only be based on an inference to the best explanation: "The objectivity of local knowledge" (as opposed to any other candidate) should be accepted on the grounds that it best explains "the possibility of planning, prediction, manipulation, control, and policy setting". The moral then is that Cartwright's recent, more robust, realism can only be based on the very method that she has taken pains to disarm. We can now move on to look at the credentials of one her stronger early arguments against IBE, viz., the alleged falsity of laws.

False Laws?

One of Cartwright's main theses is that explanation and truth pull apart (Cartwright 1983). When laws come into the picture, this thesis seems to be the outcome of a certain failure of laws. She puts it like this:

> For the fundamental laws of physics do not describe true facts about reality. Rendered as descriptions of facts, they are false; amended to be true, they lose their fundamental, explanatory power. (Cartwright 1983: 54)

So, we are invited to see that if laws explain, they are not true, and if they are true, they do not explain. What Cartwright has in mind, of course, is what she calls fundamental or abstract laws as well as the covering-law model of explanation. If laws explain by "covering" the facts to be explained, then, Cartwright says, the explanation offered will be false. If, she would go on, the laws are amended by using several *ceteris paribus* clauses, they become truer but do not "cover" the facts anymore; hence, in either case, they do *not* explain the facts. The reason why covering laws do not explain has mostly to do with the fact that the actual phenomena are too complex to be covered by simple laws. Recall her example of a charged particle that moves under the influence of two forces: the force of gravity and Coulomb's force. Taken in isolation, neither of the two laws (i.e. Newton's inverse-square law and Coulomb's law) can describe the actual motion of the charged particle. From this, Cartwright concludes that each loses either its truth or its explanatory power. Here is her argument:

> The effect that occurs is not an effect dictated by any one of the two laws separately. In order to be true in the composite case, the law must describe one effect (the effect that actually happens); but to be explanatory, it must describe another. There is a trade-off here between truth and explanatory power. (Cartwright 1983: 59)

I fail to see how all this follows. For one, it does *not* follow that there is not (worse, there cannot be) a complex law that governs the motion of massive *and* charged particles. If we keep our eyes not on epistemology (can this law be known or stated?) but on metaphysics (can there be such a law?), the above argument is, to say the least, inconclusive. For another, in the composite case, there is no formal tension between truth and explanation. In the composite case, none of the two laws (Newton's and Coulomb's) is strictly true of, in the sense of "covering", the effect that actually happens. Why should we expect each of them on its own to "cover" the complex effect? After all, the complex phenomenon is governed by both of them jointly, and hence it cannot be covered by each of them separately. This does not imply that laws lose their explanatory power. They still explain how the particle would behave if it was just massive and not charged or if it was charged but not massive. And they still contribute to the full explanation of the complex effect (that is, of the motion of the charged and massive particle). To demand of each of them to be explanatory in the sense that each of them should somehow cover the actual complex effect is to demand of them something they cannot do. The laws do not thereby cease to be true, nor explanatory. Nor does it follow that they don't jointly govern the complex effect. Governing should not be conflated with covering.[7]

My argument so far might be inconclusive. So I want to suggest that there is an important independent reason why we should take laws seriously. Laws individuate properties: Properties are what they are because of the laws they participate in. Cartwright says:

> What I invoke in completing such an explanation are not fundamental laws of nature, but rather properties of electrons and positrons, and highly complex, highly specific claims about just what behaviour they lead to in just this situation. (Cartwright 1983: 92)

If it is the case that no laws then no properties, or if properties and laws are so intertwined that one cannot specify the former without the latter, then some laws had better be true. For if they are not, then we cannot talk of properties either.[8]

This last point, however, is controversial, especially as of late. It relies on a Humean understanding of properties. And Cartwright is a non-Humean, more or less about everything. This observation is crucial. For it is Humeanism that is Cartwright's real opponent. Her capacities are non-Humean tendencies: causal powers. That is, they are irreducible, primary and causally active constituents of the world. Similarly, her properties are non-Humean properties: They are active causal agents, which are identified via their causal role and their powers. So it is not laws that determine what they are; rather, it is properties (capacities, etc.) that determine what, if any, laws hold in the world. With all this in mind, let us turn our attention to her views about

capacities. This is just one of her non-Humean themes. But it is perhaps the most central one.

CAPACITIES

Cartwright has devoted two books in the defence of the claim that capacities are prior to laws (Cartwright 1989; 1999). As is well known, she challenges the Humean view that laws are exceptionless regularities, since, she says, there are no such things.[9] How then does it appear that there *are* regularities in nature, e.g., that all planets move in ellipses?

Nomological Machines

Cartwright does not deny that there can be regular behaviour in nature. But she claims that where there is regular behaviour in nature, there is a nomological machine that makes it possible. A "nomological machine" is

> a fixed (enough) arrangement of components, or factors, with stable (enough) capacities that in the right sort of stable (enough) environment will, with repeated operation, give rise to the kind of regular behaviour that we represent in our scientific laws. (Cartwright 1999: 50)

Nomological machines make sure that "all other things are equal". So, they secure the absence of factors, which, were they present, would block the manifestation of a regularity. Take Kepler's law that all planets move in ellipses. This is not a strictly universal and unconditional law. Planets do (approximately) describe ellipses, if we neglect the gravitational pull that is exerted upon them by the other planets, as well as by other bodies in the universe. So, the proper formulation of the law, Cartwright argues, is: *ceteris paribus*, all planets move in ellipses. Now, suppose that the planetary system is a stable enough nomological machine. Suppose, in particular, that as a matter of fact, the planetary system is (for all practical purposes) shielded: It is sufficiently isolated from other bodies in the universe, and the pull that the planets exert on each other is negligible. Under these circumstances, we can leave behind the *ceteris paribus* clause and simply say that all planets move in ellipses. But the regularity holds only so long as the nomological machine backs it up. If the nomological machine were to fail, so would the regularity. As Cartwright has put it: '(L)aws of nature (in this necessary regular association sense of "law") hold only *ceteris paribus*—they hold only relative to the successful repeated operation of a nomological machine' (Cartwright 1999: 49–50).

Nomological machines might occur naturally in nature. The planetary system is such a natural nomological machine. But, according to Cartwright, this is exceptional. As she says: 'more often [the nomological machines] are

engineered by us, as in a laboratory experiment' (Cartwright 1999: 49). 'In any case', she adds, 'it takes what I call a nomological machine to get a law of nature' (Cartwright 1999: 49).

For the operation of a nomological machine, it is not enough to have a stable (and shielded) arrangement of components in place. It is not enough, for instance, to have the sun, the planets, and the gravitational force in place in order for the planetary machine to run. Cartwright insists that it is the *capacities* that the components of the machine have that generate regular behaviour. For instance, 'a force has the capacity to change the state of motion of a massive body' (Cartwright 1999: 51). Couldn't the nomological machine itself be taken to be a regularity? No, she answers: 'the point is that the fundamental facts about nature that ensure that regularities can obtain are not again themselves regularities. They are facts about what things can do' (Cartwright 1995: 4). But what exactly are capacities, i.e., the things that things can do?

Cartwright focused her attention on 'what capacities do and why we need them' and *not* on 'what capacities are' (Cartwright 1989: 9). What they are is the job of her *The Dappled World*. Before, however, we examine what they are, let us see the main argument she offers as to why we need capacities.

Why Do We Need Capacities?

The Sellarsian Argument

Sellars's master argument for commitment to the unobservable entities posited by scientific theories is that they play an ineliminable explanatory role (Sellars 1963). In order to formulate it, he had to resist what he aptly called the 'picture of the levels'. According to this picture, the realm of facts is layered. There is the bottom level of observable entities. Then, there is an intermediate (observational) level of empirical generalisations about observable entities. And finally, there is yet another (higher-theoretical) level: unobservable entities and laws about them. It is part of this picture that while the observational framework is explanatory of observable entities, the theoretical framework enters the picture by explaining the inductively established generalisations of the observational framework. But then, Sellars says, an empiricist can protest that the higher level is dispensable. He may argue that all the explanatory work vis-à-vis the bottom level is done by the observational framework and its inductive generalisations. Why then, he may wonder, posit a higher level in the first place?

Sellars's diagnosis is that this picture rests on a myth. His argument against the myth of the levels is that the unobservables posited by a theory explain directly why (the individual) observable entities behave the way they do and obey the empirical laws they do (to the extent that they do obey such laws). So, he resists the idea that the theoretical framework has

as its prime function to explain the empirical generalisations of the observational framework. Sellars claimed that unobservable entities are indispensable because they also explain why observational generalisations are, occasionally, violated; why, that is, some observable entities do not behave they way they should, had their behaviour been governed by the observational generalisation.

This is a fine argument and I endorse it fully (Psillos 2004a). Cartwright offers an argument structurally similar to Sellars's in defence of capacities (Cartwright 1989: 163). She has in mind another possible layer cake. The bottom level is the nonmodal level of occurrent regularities; the intermediate level is the level of Humean laws (either deterministic or statistical). The higher level is supposed to be a *sui generis* causal one. *This* layer cake, Cartwright notes, also invites the thought (or the temptation) to do away with the higher level altogether. All the explanatory work, it might be said, is done by Humean laws, endowed with modal force. The higher (causal) level could then be just seen as a higher modal level, with no claim to independent existence: It is just a way to talk about the intermediate level, and in particular a way to set constraints on laws in order to ensure that they have the required modal force. It is *this* layer cake that Cartwright wants to resist. For her, the higher causal level is indispensable for the explanation of what regularities there are (if any) in the world. So we seem to have a solid Sellarsian argument for capacities. But do we?

Capacities and Regularities

Before we proceed to examine this, an exegetical point is in order. Cartwright splits the higher (causal) level into two sublevels: a lower sublevel of *causal laws* and a higher sublevel of ascriptions of capacity. She couches all this in terms of two levels of generality or more accurately of two levels of modality (Cartwright 1989: 142). She says:

> (. . .) the concept of general *sui generis* causal truths—general causal truths not reducible to associations—separates naturally into two distinct concepts, one at a far higher level of generality than the other: at the lower level we have the concept of a causal law; at the higher, the concept of capacity. I speak of two levels of generality, but it would be more accurate to speak of levels of modality, and for all the conventional reasons: the claims at both levels are supposed to be universal in space and through time, they support counterfactuals, license inferences, and so forth. (Cartwright 1989: 142)

Why do we need *two* causal levels? Why, in particular, do we need a level of capacities? To cut a long story short, Cartwright thinks that causal laws are kinds of causal generalisations relative to a particular population

(Cartwright 1989: 144). They are causal, as opposed to Humean laws of association, mostly because, as Cartwright argues, the facts they report (e.g., that aspirins relieve headaches or that smoking causes cancer) cannot be fully captured by probabilistic relations among magnitudes or properties. Causal information is also required to specify the conditions under which they hold. A further thought then is that ascription of capacities is also necessary in order to remove the relativised-to-a-population character of causal laws. We don't just say that smoking causes cancer to population *X*. We also want to say that smoking causes cancer, *simpliciter*. This claim (which is universal in character) is best seen as a claim about capacities: *C* causes *E* means *C* carries the capacity *Q* to produce *E* (Cartwright 1989: 145). Capacities, then, are introduced to explain causal laws and to render them universal in character.[10] This last point is crucial: Causal laws are *ceteris paribus*. After all, it's not invariably the case that aspirin relieves headache. But capacities remove the *ceteris paribus* clause: Aspirin *always* carries the capacity to relieve headache. Capacities, we are told, are *stable*. If something has the capacity *Q*, then it carries it with it from one situation to another (Cartwright 1989: 145).

What then of Cartwright's Sellarsian argument for capacities? I focus on just one central problem. Sellars saves the higher level of electrons, protons, etc. by focusing on the indispensable role this level plays in the explanation of singular observable phenomena or things. Similarly, one would demand of Cartwright's argument to show how capacities are indispensable for the explanation of occurrent regularities, without the intervening framework of Humean laws plus modal force. But it seems that there is a tension in her argument. Whereas in Sellars's case, the entities of the theoretical framework (unobservables) can be identified independently of the entities in the bottom framework, it is debatable that this can happen in Cartwright's case. Here there are conflicting intuitions. One is that we need regularities (or Humean laws) to identify what capacities things carry. Another (Cartwright's, I think) is that this is not the case. I am not entirely certain whose intuitions are right. But it seems to me that the Humean is on a better footing. Capacities might well be posited, but only after there has been a regular association between relevant event types. No one would mind ascribing to aspirin the capacity to relieve headaches, if that was the product (as indeed it is) of a regular association between taking aspirins and headaches going away. "Regular" here does not necessarily mean exceptionless. But, so much the better for positing capacities if the association happens to be exceptionless. To say the least, one could more easily explain how capacities have modal force. So, there is an important disanalogy between Sellars's argument for unobservables and Cartwright's argument for capacities, which casts doubt on the indispensability of positing capacities. That is, in Cartwright's case, we need the lower level (regularities) to identify the entities of the higher level (capacities).

Single Cases

Cartwright insists that capacities might reveal themselves only occasionally or only in a single case. Consider what she says:

> "Aspirins relieve headaches". This does not say that aspirins always relieve headaches, or always do so if the rest of the world is arranged in a particularly felicitous way, or that they relieve headaches most of the time, or more often than not. Rather it says that aspirins have the capacity to relieve headaches, a relatively enduring and stable capacity that they carry with them from situation to situation; a capacity which may if circumstances are right reveal itself by producing a regularity, but which is just as surely seen in one *good* single case. The best sign that aspirins can relieve headaches is that on occasion some of them do. (Cartwright 1989: 3, emphasis added)

This is surely puzzling. Just adding the adjective "good" before the "single case" does not help much. A "good" controlled experiment might persuade the scientist that he has probably identified some causal agent. But surely, commitment to it follows only if the causal agent has a regular behaviour that can be probed in similar experiments. A single finding is no more compelling than a single sighting of a UFO. Single or occasional manifestations cast doubt on the claim that there is a stable and enduring capacity at play (Glennan 1997: 607–608).

Cartwright disagrees. She advances what she calls the "analytic method" in virtue of which capacity ascriptions are established (Cartwright 1999) and later summarises her ideas thus:

> We commonly use the analytic method in science. We perform an experiment in "ideal" conditions, I, to uncover the "natural" effect E of some quantity, Q. We then suppose that Q will in some sense "tend" or "try" to produce the same effect in other very different kinds of circumstances. (. . .) This procedure is not justified by the regularity law we establish in the experiment, namely 'In I, $Q \rightarrow E$'; rather, to adopt the procedure is to commit oneself to the claim "Q has the capacity to E". (Cartwright 2002: 435–436)

What is the force of this claim? Note, first, that we don't have a clear idea of what it means to say that Q "tends" or "tries" to produce its effects. It seems that either Q does produce its effect or it doesn't (if, say, other factors intervene). Second, as Teller notes, it is not clear how the "trying" can be established by looking at a single case only (Teller 2002: 718). One thought here might be that if we have seen Q producing its effect at least one time, we can assume that it can produce it; and hence that it has the capacity to

produce it. But I don't think this is the right way to view things. Consider the following three questions: (i) what exactly is *Q*'s effect? (ii) how can we know that it was *Q* which brought E about? and (iii) wouldn't it be rather trivial to say that for each effect there is some capacity *X* which produces it? All three questions would be (more easily) answered if we took capacities to be regularly manifested. The "regularity law", "in *I*, $Q \rightarrow E$" makes the positing of a capacity legitimate. It is because (and insofar as) "in *I*, $Q \rightarrow E$" holds that we can say that "*Q* has the capacity to *E*" and not the other way around.[11]

If the capacity *Q* of *x* to bring about *y* was manifested regularly, then one could say that the presence of the capacity could be tested. Hence, one could move on to legitimately attribute this capacity to *x*. But if a capacity can manifest itself in a single case, it is not clear at all how the presence of the capacity can be tested. Why, in other words, should we attribute to *x* the capacity to bring about *y*, instead of claiming that the occurrence of *y* was a matter of chance? So, there seems to be a tension between Cartwright's claim that capacities are manifestable even in single cases and her further claim that capacities are testable.[12]

So far, I have focused on the relation between capacities (the higher level) and regularities (the lower level). But there is also a problem concerning the two sublevels of the higher level, viz., capacities and causal laws.[13] Do claims about the presence of capacities have extra content over the claims made by ordinary causal laws? So, do we really need to posit capacities? Take, for instance, the ordinary causal law that aspirin relieves headaches. If we ascribed to aspirin a capacity to relieve headaches, would we gain in content? There is a sense in which we would. Ordinary causal laws are *ceteris paribus*, whereas capacity claims are not. Because it is only under certain circumstances that aspirin relieves headaches, it is only *ceteris paribus* true that aspirin causes headache relief. But, Cartwright might say, once it is established that aspirin carries the capacity to relieve headaches, the *ceteris paribus* clause is removed: The capacity is always there, even if there may be contravening factors that block, on occasion, its manifestation. The problem with this attempt to introduce capacities is that the strictly universal character of claims about capacities cannot be established. If it is allowed that claims about the presence of capacities might be based on single manifestations, it is not quite clear what kind of inference is involved in the movement from a single manifestation to the presence of the capacity. Surely, it cannot be an inference based on any kind of ordinary inductive argument.[14] If, on the other hand, it is said that claims about capacities are established by ordinary inductive methods, based on several manifestations of the relevant capacity, then all that can be established is a *ceteris paribus* law. Based on cases of uses of aspirin, all that it can be established is that *ceteris paribus*, aspirin relieves headaches. So, it is questionable that talk about capacities has extra content over talk about ordinary causal laws.

Cartwright could argue that claims about capacities are strictly universal in the sense that objects have capacities even if they completely fail to manifest them (Cartwright 2002: 427–428). However, she would then seem to compromise her view that capacities are measurable and testable. There is a deep, if common, reason why we should be wary of unmanifestable capacities: There could be just too many of them, even contradictory ones. Couldn't we just say of any false generalisation (e.g., that bodies rise if they are left unsupported) that the bodies referred to in it have the relevant capacity, though it is never manifested? And couldn't we say that an object carries at the same time the stable capacity to rise if left unsupported and the stable capacity to fall if left unsupported, but that the former is unmanifestable? In other words, what distinguishes between unmanifestable capacities and nonexistent ones?

Moral: if Cartwright insists on single manifestation of capacities, she faces a sticky trilemma. Either talk of capacities does not have extra content over talk in terms of ordinary causal laws; or there is a mysterious method that takes us from a single manifestation to the capacity; or there are unmanifestable capacities. All three options have unpalatable consequences.

Capacities and Interactions. To be fair to Cartwright, she has offered other reasons for commitment to capacities. One of them is that capacities can explain causal interaction. She says that 'causal interactions are interactions of causal capacities, and they cannot be picked out unless capacities themselves are recognised' (Cartwright 1989: 164).

There are cases that fit this model. A venomous snake bites me, and I take an antidote. The venom in my bloodstream has the capacity to kill me, but I don't die because the antidote has the capacity to neutralise the venom. That's a case of causal interaction, where one capacity blocks another. I am not sure this commits us to *sui generis* capacities, as opposed to whatever chemical properties the venom and the antidote have, and a law that connects these properties. But let's not worry about this. For there is a more pressing problem.

Suppose that I take an aspirin while I am still hearing the continuous and desperate screaming of my daughter, who suffers from colic. The aspirin has the capacity to relieve my headache, but the headache does not go away. It persists undiminished. How shall I explain this? Shall I say that this is because the screaming of my daughter has the capacity to cause aspirin-resistant headaches? This would be overly ad hoc. Shall I say that this is because the screaming of my daughter has the capacity to neutralise the capacity of aspirin to relieve headache? This would be very mysterious. Something has indeed happened: There has been an interaction of some sort which made aspirin not work. But why should I attribute this to a capacity of the screaming? If I did that, I would have to attribute to the screaming a number of capacities: the capacity to-let-aspirin-work-if-it-is-mild, the capacity to let-aspirin-work-if-it-is-*not*-mild-but-I-go-away-and-let-my-wife-deal-with-my-daughter, the capacity to block-aspirins'-work-if-it-is-extreme, etc. This

is not exactly an argument against the role of capacities in causal interaction (though it might show that there can be causal interaction without reference to capacities). Still, it seems a genuine worry: When trying to account for causal interaction, where do we stop positing capacities and what kinds should we posit?

Cartwright challenges the sceptic about capacities with the following: 'the attempt to "modalise away" the capacities requires some independent characterisation of interactions; and there is no general non-circular account available to do the job' (Cartwright 1989: 164). If we could not characterise interactions without reference to capacities, we had better accept them. But why not follow, for instance, Salmon (1997) or Dowe (2000) in their thoughts that interactions are explained in terms of exchanges of conserved quantities? There is no compelling reason to take *them* to be capacities. We could; (Cartwright, for instance, takes charge to be a capacity). But then again we couldn't. Charge might well be a property (an occurrent property, that is) in virtue of which, and given certain laws, a particle that instantiates it behaves the ways it does.[15]

What Are Capacities?

Suppose that we do need to posit capacities. What exactly is the thing we need to posit? Cartwright is certainly in need of a more detailed account of how capacities are individuated. She tells us that capacities are *of* properties and not *of* individuals: 'the property of being an aspirin carries with it the capacity to cure headaches' (Cartwright 1989: 141). But aspirin is not, strictly speaking, a property. It's something that has a property. And certainly it does not carry its capacity to relieve headaches in the same way in which it carries its shape or colour.

It would be more accurate to say that capacities are properties of properties. That is, that they are second-order properties. But this would create some interesting problems. It would open the way for someone to argue that capacities are functional (or causal) roles. This would imply that there must be occupants of these causal roles, which are not themselves capacities. They could be the properties (maybe many and variable) that occupy this causal role. So, the capacity to relieve pain would be a causal role filled (or realised) by different properties (e.g., the chemical structure of paracetamol or whatever else). If, however, we take capacities to be causal roles, it would be open for someone to argue, along the lines of Prior, Pargeter, and Jackson (Prior et al. 1982) that capacities are causally impotent. The argument is simple. Capacities are distinct from their causal bases (as they are properties *of* them). They must have a causal basis (a realiser) because they are second-order. This causal basis (some properties) are themselves a sufficient set of properties for the causal explanation of the manifestation of the capacity (whenever it is manifested). Hence, the capacity *qua* distinct (second-order) property is causally impotent.

Cartwright wouldn't be willing to accept this conclusion. But then capacities must be *of* properties (or be carried by properties) in a different way. What exactly this way is it is not clear. She asks: 'Does this mean that there are not one but two properties, with the capacity sitting on the shoulder of the property which carries it?' And she answers: 'Surely not' (Cartwright 1989: 9). But no clear picture emerges as to what this relation of "*a* carrying *b*" is. (And is this "carrying" another capacity, as in *a* has the capacity to carry *b*? And if so, isn't there a regress in the offing?) At a different place, we are told that capacities have powers, which they can retain or lose (in causal interactions; Cartwright 1989: 163). Is that then a third-order property? A property (power) of a property (capacity) of a property (aspirin)? I don't think Cartwright wants to argue this. But what does she want to argue?

Cartwright later returns to these issues (Cartwright 1999). Here it seems that another possibility is canvassed, viz., that properties themselves are capacities. It's not clear whether she takes all properties to be capacities, but it seems that she takes at least some to be. We are given examples such as force and charge. I am not sure I have this right, but it seems to follow from expressions such as: 'Coulomb's law describes a capacity that a body has qua charged' (Cartwright 1999: 53). It also seems to follow from considering concepts such as 'attraction, repulsion, resistance, pressure, stress, and so on' as concepts referring to capacities (Cartwright 1999: 66). In any case, it seems that she aligns herself with Shoemaker's view of properties as "conglomerates of powers" (see Cartwright 1999: 70). Capacities then seem to come more or less for free: 'Any world with the same properties as ours would ipso facto have capacities in it, since what a property empowers an object to do is part of what it is to be that property' (Cartwright 1999: 70). So, it seems that Cartwright adopts a causal theory of properties, where properties themselves are causal powers.

Capacities and Laws

A number of questions crop up at this point. First, are all powers with which a property empowers an object constitutive of this property? And if not, how are we to draw a distinction between constitutive powers and nonconstitutive ones? For instance, is the causal power of aspirin to relieve headache on a par with its causal power to produce a pleasing white image? This is not a rhetorical question. For it seems that in order to distinguish these two powers in terms of their causal relevance to something being an aspirin, we need to differentiate between those powers that are causally relevant to a certain effect, e.g., relieving pain, and those powers that are not. Then, we seem to run in the following circle. We need to specify what powers are causally relevant to something being *P*. For this, we need to distinguish the effects which are brought about by *P* in two sorts: those that are the products of causally relevant powers and those that are not. But in order to do this we need first to specify what it *is* for something to be *P*.[16] That is, we need to specify what

powers are causally relevant to *P*'s identity and what are not. Ergo, we come back to where we started. (Recall that on the account presently discussed causal powers are the *only* vehicle to specify *P*'s identity).

Second question: why is it the case that some causal powers are held together, and others are not? Why, that is, do certain powers have a certain kind of "causal unity", as Shoemaker (1980: 125) put it? This is a crucial question because even if every property is a cluster of powers, the converse does not hold. Electrons come with the power to attract positively charged particles and the power to resist acceleration, but they don't come with the power to be circular. And the power of a knife to cut wood does not come with the power to be made of paper. This is important because, as Shoemaker himself observes, the concurrence of certain powers might well be the consequence of a law (1980: 125). So, it might well be that laws hold some capacities together. Hence, it seems that we cannot just do with capacities. We also need laws as our building blocks. This issue has a ramification. Why is it the case that nothing has the power to move faster than light? The absence of a certain capacity might also be the consequence of a natural law.

Third question: should we be egalitarian about capacities? Is the capacity to resist acceleration on a par with the capacity to become grandparent? Or with the capacity to be a table-owned-by-George-Washington? This question is different from the first. It relates to what in the literature is called the difference between genuine changes and mere Cambridge changes. The parallel here would be a difference between genuine capacities (properties) and mere Cambridge capacities (properties). Here again, laws are in the offing. For it can be argued that genuine capacities (properties) are those that feature in laws of nature.

I offer these questions as challenges. But they do seem to point to a certain double conclusion. On the one hand, we need to be told more about what capacities are before we start thinking seriously that we should be committed to them. On the other hand, we seem to require laws as well as capacities, even if we accept capacities as building blocks.

Cartwright wants to further advance the view that capacities are metaphysically prior to laws. She says, 'It is capacities that are basic, and laws of nature obtain—to the extent that they do obtain—on account of the capacities' (Cartwright 1999: 49). She offers no formal treatment of the issue how capacities relate to laws. Instead, we are given some examples.

> I say that Newton's and Coulomb's principles describe the capacities to be moved and to produce a motion that a charged particle has, in the first case the capacity it has on account of its gravitational mass and in the second, on account of its charge. (Cartwright 1999: 65)

If laws describe what the entities governed by them can do on account of their capacities, these capacities should be individuated, and ascribed, to

entities, independently of the law-like behaviour of the latter. But, as noted above, it is not clear that this can be done. It seems that far from being independent of laws, the property of, say, charge is posited and individuated by reference to the law-like behaviour of certain types of objects: Some attract each other, and others repel each other in a regular fashion. The former are said to have opposite charges, whereas the latter have a similar charge. Cartwright says: 'The capacity is associated with a single feature—charge—which can be ascribed to a body for a variety of reasons independent of its display of the capacity described in the related law' (Cartwright 1999, 54–55).

This may well be true. But it does not follow that the capacity is grounded in no laws at all. Cartwright disagrees. She argues that '[c]apacity claims, about charge, say, are made true by facts about what it is in the nature of an object to do by virtue of being charged' (Cartwright 1999: 72).

Then, one would expect an informative account of what it is in the nature of an object to do. Specifically, one would expect that the nature of an object would determine its capacities, and would delineate what this object can and cannot do. But she goes on to say: 'There is no fact of the matter about what a system can do just by virtue of having a given capacity. What it does depends on its setting . . . (Cartwright 1999: 73).

Why, then, should we bother to attribute capacities? We could just offer an open-ended list of the things that a system does when it is placed in several settings. If, at least, there was a fact of the matter as to what a system could do by virtue of having a given capacity, the capacity could be used to (a) predict what a system can or cannot do and (b) explain why it behaves the way it does. In fact, if Cartwright really means to uphold the strong view that there is no fact of the matter as to what a system can do by having a certain capacity, then the very possibility of prediction and of explanation is threatened. For any kind of behaviour would be compatible with the system's having a certain capacity. No specific behaviour could be predicted, and any kind of behaviour could be explained (by an appeal to context-specific impediments of the system's capacities).[17]

One might object, however, that Cartwright's wording is very careful. It does not imply that there is no fact of the matter about what a system (or an object) can do by virtue of its nature. Yet, one would expect that if the nature of an object placed some substantive constraints on its capacities, there would be a fact of the matter about what this object can do by virtue of its capacities. For instance, one would expect that although a certain particle has the capacity to move, its nature constrains this capacity so that it cannot move with velocity greater than the velocity of light. As this example suggests, it might well be the case that the nature of an object is constrained by what laws it obeys.

In a previous draft of this paper, I tried to examine in some detail what these natures are and how they might relate to capacities. But Paul Teller directed my attention to the following passage, in which Cartwright says:

> My use of the terms *capacity* and *nature* are closely related. When we ascribe to a feature (like charge) a generic capacity (like the Coulomb capacity) by mentioning some canonical behaviour that systems with this capacity would display in ideal circumstances, then I say that that behaviour is *in the nature of* that feature. Most of my arguments about capacities could have been put in terms of natures. . . .
>
> (Cartwright 1999: 84–85)

So, it seems clear that Cartwright thinks there is no significant distinction between capacity and nature. But suppose that she followed many other friends of capacities and distinguished between capacities and natures. Fisk (1970) and Harré (1970), among others, think that an appeal to an entity's nature can explain why this entity has certain capacities. In particular, Harré (1970) argues that (a) discovering the nature of an entity is a matter of empirical investigation; but (b) specifying (or knowing) the exact nature of an entity is not necessary for grounding the ascription of a power to it. He links natures and capacities thus: 'There is a ϕ such that something has φ, and whatever had φ in *C, would have* to *G*, i.e. if something like α did not have φ in *C* it would not, indeed *could* not *G* (Harre 1970: 101).

The nature φ of an entity is thereby linked with its capacity to *G*. There are many problems with this proposal.[18] But I focus on one. What is it that makes the foregoing counterfactual true? It's not enough to have the circumstances *C* and the nature φ in order to get *G*. This is not just because *G* could be unmanifested. Even if we thought that the power to *G* were always manifested in circumstances *C* with a characteristic effect *e*, there would *still* be room for asking the question: What makes it the case that α's being φ in *C* makes it produce the characteristic effect *e*? We need, that is, something to relate (or connect) all these together, and the answer that springs to mind is that it is a *law* that does the trick.[19] This law might well be a brute (Humean) regularity.[20]

An advocate of natures could say that when the nature φ is present, there is no need to posit a law in order to explain why a certain object has a characteristic effect *e* when the circumstances are *C*. Yet this move would not really be explanatory. It would amount to taking natures to be collections of powers, and this hardly explains in an interesting way why a certain nature has the capacities it does: It just equates the nature of an object with a collection of its capacities.

A CONCLUDING REMARK

As we have seen, Cartwright has moved from a modest realist position (viz., realism about entities) to a superrealist position (viz., realism about powers and capacities). Part of her motivation for her early, restricted, realism was

a certain antifundamentalism, viz., a resistance to the view that there are fundamental laws of nature, which determine what entities do, and which are captured (or should be captured) by scientific theories. It may be ironic that she now replaces this picture by another fundamentalism, viz., the view that capacities are the fundamental building blocks of the world, the things that make things to be what they are and to behave the way they do. Along the way, her early antitheory temper was softened. But her early antilaws temper was hardened.

In contemplating Cartwright's realist toil, we have learned a lot. But it seems that we are still short of a compelling reason to take capacities seriously as fundamental non-Humean constituents of the world. At any rate, even if we granted capacities, we would still need laws to (i) identify them; (ii) connect them with their manifestations; (iii) explain their stability; (iv) explain why some (but not others) occur together; (v) explain why some (but not others) obstruct the manifestation of others. It seems then that both the epistemology and the metaphysics of capacities require laws. Cartwright is to be commended for trying to make a case for the view that *capacities are enough for laws*. If the argument in the later part of this paper has been correct, then the situation is more complicated: *Laws and capacities are necessary for laws*.

NOTES

1. Earlier versions of this chapter were presented at the Workshop in honour of Nancy Cartwright, in Konstanz, December 2002, and in the University of California San Diego Philosophy Colloquium. I thank the participants of these events for many thoughtful comments and criticisms. I especially thank Nancy Cartwright for her comments and encouragement, as well as Craig Callender, Paul Churchland, Gerald Doppelt, Ron Giere, Stephan Hartmann, Carl Hoefer, and Wolfgang Spohn. Paul Teller deserves special mention for giving me many thoughtful written comments on the content as well as the structure of this chapter. Theodore Arabatzis, Steve Clarke, Robin Hendry, Christoff Schmidt-Petri, and David Spurrett must also be thanked for detailed written comments.
2. For more on the relation between scientific realism and metaphysical issues, see Psillos 2005.
3. For readers unfamiliar with these attempts, a brief statement of some major views follows. On Lewis's reading, *c* causally explains *e* if *c* is connected to *e* with a network of causal chains. For him, causal explanation consists in presenting portions of explanatory information captured by the causal network. On Woodward's reading, *c* causally explains *e* if *c* and *e* are connected by a relevant (interventionist) counterfactual of the form 'if *c* hadn't happened, then *e* wouldn't have happened either'. On Salmon's reading, *c* causally explains *e* if *c* is connected with *e* by a suitable continuous causal (i.e. capable of transmitting a mark) process. On the standard deductive-nomological reading of causal explanation, for *c* to causally explain *e*, *c* must be a nomologically sufficient condition for *e*. And for Mackie, for *c* to causally explain *e* there must be event-types *C* and *E* such that *C* is an inus-condition for *E*. For details on all these, see Psillos (2002a).

4. For a different take on the nature of inference to the most likely cause, see Suárez's contribution in this volume.
5. For more on this issue see Psillos (1999: Ch. 4).
6. To see what these worries might be, consider the difference between modest and ambitious transcendental arguments. Is Cartwright's intention to arrive at the modest conclusion that it is rational to believe that there is local knowledge or at the much more ambitious conclusion that *there is* local knowledge?
7. Spurrett defends a similar point in much more detail (Spurrett 2001).
8. A huge issue here concerns the nature of laws. I favour the Mill–Ramsey–Lewis approach, which I defend in some detail in Psillos (2002a: 148–154, 210–211). This approach can identify laws independently of their ability to support counterfactuals. However, it seems to require some prior notion of 'natural property'. But this notion need not equate properties with causal powers or capacities.
9. For an important survey of the debate around *ceteris paribus* laws, as well as a defence of strict laws in physics, see Earman& Roberts (1999). The interested reader should also see the special issue of *Erkenntnis* (2002, Vol. 57, no 3) on the status of *ceteris paribus* laws.
10. This point is also made vividly in Cartwright (2002).
11. Cartwright argues that capacities help explain what makes the design of a single experiment 'a good one': The design is good if it controls for all factors relevant to the effect (Cartwright 2002: 436). But why do we need an appeal to capacities to do this? In a clinical trial what Cartwright demands can be achieved by randomisation. In a physical experiment, in order to control for all factors relevant to the effect we need to appeal to regularities in the following sense: we need to control for all factors that regularly influence effects of this type. Strictly speaking, we cannot control for factors that do not fall under a regularity, since we don't have a clue as to what they might be. When, in an experiment, one does not control for the colour of the experimenter's eyes, it is because there is no regularity that connects the colour of eyes with the result of the experiment. Little (if anything) is gained if we add that the colour of the eyes does not have the *capacity* to alter the effect.
12. One might argue that there are clear cases in which a single case is enough to posit a capacity. An example put to me by Christoph Schmidt-Petri is the capacity to run fast: One case is supposed to prove its existence. I am not so sure this is true. What if I run fast (just once) because I took a certain steroid on a given occasion? Surely, in this case I don't have the capacity to run fast, though the steroid might have the (stable) capacity to make people run fast. But this latter capacity would need regular manifestation in order to be posited. For more criticism of Cartwright's argument that capacities are necessary in the methodology of science, see Giere (this volume).
13. A variant of this problem has been posed by Morrison (1995)
14. This point calls into question Cartwright's claim that capacities show how we can make sense of inductions from single experiments (Cartwright 1999: 90; 2002: 436). Undoubtedly, *if* stable capacities are in operation, then knowing them is enough to generalise from a single experiment. But how is the antecedent of the conditional grounded? It seems that we need regular behaviour (and hence plenty of inductive evidence) in order to posit stable capacities in the first place.
15. For a criticism of causal powers, see Psillos (2006).
16. Compare: something could be an aspirin without having the causal power to produce a white image; but something could *not* be an aspirin without having the power to relieve headaches.

17. A similar complaint is voiced by Earman & Roberts (1999: 456) and Teller (2002: 719).
18. See the criticisms of Fisk's views by Aune (1970) and McMullin (1970)
19. A similar point is made by Menzies (2002). Teller also notes that capacities might well be no different from the OK properties that Cartwright argues should figure in laws (Teller 2002: 720–721).
20. This is just one option, of course; see also Teller (2002: 722). Another option would be to look for a mechanism that connects the nature φ with its power to produce a characteristic effect in certain circumstances. I have a number of objections to mechanisms that I cannot repeat here (see Psillos 2004b). At any rate, it seems enough for the purposes of this chapter that it remains an open option that Humean regularities may get the capacities do whatever they do.

REFERENCES

Aune, B. (1970) 'Fisk on capacities and natures', *Boston Studies in the Philosophy of Science*, 8: 83–87.
Cartwright, N. (1983) *How the Laws of Physics Lie*, Oxford: Clarendon Press.
———. (1989) *Nature's Capacities and Their Measurement*, Oxford: Clarendon Press.
———. (1995) 'Précis on *Nature's Capacities and Their Measurement*", *Philosophy and Phenomenological Research 55*: 153–6.
———. (1999) *The Dappled World*, Cambridge: Cambridge University Press.
———. (2002) 'In favour of laws that are not *ceteris paribus* after all', *Erkenntnis*, 57: 425–439.
Dowe, P. (2000) *Physical Causation*, Cambridge: Cambridge University Press.
Earman, J., and J. Roberts. (1999) '*Ceteris Paribus*, there is no problem of provisos', *Synthese*, 118: 439–478.
Fisk, M. (1970) 'Capacities and natures', *Boston Studies in the Philosophy of Science*, 8: 49–62.
Giere, R. (this volume) 'Models, metaphysics and methodology'.
Glennan, S. (1997) 'Capacities, universality, and singularity', *Philosophy of Science*, 64: 605–626.
Harré, R. (1970) 'Powers', *The British Journal for the Philosophy of Science*, 21: 81–101.
Hempel, C. G. (1965) *Aspects of Scientific Explanation*, New York: The Free Press.
Hitchcock, C. R. (1992) 'Causal explanation and scientific realism', *Erkenntnis*, 37: 151–178.
Lipton, P. (1991) *Inference to the Best Explanation*, London: Routledge.
Menzies, P. (2002) 'Capacities, natures and pluralism: A new metaphysics for science?' *Philosophical Books*, 43: 261–270.
McMullin, E. (1970) 'Capacities and natures: An exercise in ontology', *Boston Studies in the Philosophy of Science*, 8: 63–83.
Morrison, M. (1995) 'Capacities, tendencies and the problem of singular causes', *Philosophy and Phenomenological Research*, 55: 163–168.
Prior, E., R. Pargeter, and F. Jackson. (1982) 'Three theses about dispositions' *American Philosophical Quarterly*, 19: 251–256.
Psillos, S. (1999) *Scientific Realism: How Science Tracks Truth*, London: Routledge.
———. (2002a) *Causation and Explanation*, Chesham: Acumen.
———. (2002b) 'Simply the best: A case for abduction', in F. Sadri and A. Kakas (eds) *Computational Logic: From Logic Programming into the Future*, LNAI 2408, Berlin-Heidelberg: Springer-Verlag.

———. (2004a) 'Tracking the real: Through thick and thin', *The British Journal for the Philosophy of Science*, 55: 393–409.
———. (2004b) 'A glimpse of the *Secret Connexion*: Harmonising mechanisms with counterfactuals', *Perspectives on Science*, 12: 288–319.
———. (2005) 'Scientific Realism and Metaphysics', *Ratio*, 18, 385–404—*Ratio, 18*: 385–404.
———. (2006) 'What do powers do when they are not manifested?' *Philosophy and Phenomenological Research*, 72: 135–156—'*Philosophy and Phenomenological Research,* 72: 135–156.
Salmon, W. (1997) *Causality and Explanation*, Oxford: Oxford University Press.
Sellars, W. (1963) *Science, Perception and Reality*, Atascadero: Ridgeview P. C.
Suárez, M. (this volume) 'How inference to the most probable cause might be sound'.
Shoemaker, S. (1980) 'Causality and properties', in P. van Inwagen (ed.) *Time and Change*, New York: Springer.
Spurret, D. (2001) 'Cartwright on laws of composition', *International Studies in the Philosophy of Science*, 15: 253–268.
Teller, P. (2002) 'Review of *The Dappled World*', *Nous*, 36: 699–725.

Reply to Stathis Psillos

Stathis Psillos demands a way to identify capacities. It seems we either need laws—'laws individuate properties; properties are what they are because of the laws they participate in'—or a set of behaviours that occur when the capacity is manifested (Psillos this volume: 15). But, he observes, I don't like laws, and I say that some capacities can be manifested in almost any behaviour. Neither of these claims is entirely accurate, however.

The law claims I don't like are those that report regular associations among occurrent properties. But there are other "laws" that I endorse wholeheartedly; for instance, "An object of mass m has a capacity of strength GMm/r^2 to attract a mass of size M a distance r away". This law ascribes a given capacity to a property that we have other ways to identify.[1] Or, "If an object of mass m manifests its capacity to attract an object of mass M a distance r away and nothing interferes, the second object will have an acceleration Gm/r^2." Notice that in this last case we also have a claim about what behaviour occurs when the capacity is manifested. I can thus mimic Psillos: A given capacity is what it is because of the laws it participates in. These laws often involve reference to other capacities, but that is no more an objection to the claim that the laws "individuate" the capacity than the fact that the laws that are supposed to individuate a property refer to other properties.

Some of Psillos's worries about identifying capacities by their manifestation rest on a conflation of the manifestation, or exercise, of the capacity with the occurrence of the canonical behaviour we associate with the capacity.[2] The gravitational capacity, for instance, seems always to manifest itself—a massive object always *attracts* another, yet the canonical behaviour—an acceleration towards that object—may seldom occur. And we know a host of tests that assure us that the manifestation obtains even when the acceleration does not. So the manifestation—"attracting"—is fixed even if the behaviour described in occurent-property language—"acceleration Y"—is highly various.[3]

Psillos also worries about prediction and explanation. True, for many capacities almost any behaviour can result from their exercise. But we can still predict because different behaviours result from the exercise of the same

capacity in different circumstances. So long as there are rules about how capacities combine or how they respond to variations in circumstance, prediction will be possible.

Psillos share a different worry about testing with Margaret Morrison.[4] I claim that capacities can often be measured and very precisely. They are like forces in that respect. But that does not mean that we can tell by those measurements that what we measure is a capacity. Again, that is like forces. We can measure the acceleration of an object and its mass and multiply to measure the force on it. That does not tell us that there are forces in nature. To defend this, we need an extensive network of empirical, theoretical and philosophical considerations. So too with capacities.[5]

As to what capacities are, I do not object to the putatively untoward consequences of either alternative Psillos offers. Suppose for instance that ". . . is an interference with . . .", ". . . has the capacity to . . .", ". . . is a trigger for . . .", etc., are second-order properties. What matters for capacities is the threefold distinction Hume denied between the obtaining of the capacity (e.g., the capacity to attract obtains whenever an object has a mass), the manifestation or exercise of the capacity (the attracting),[6] and the "occurent-property" behaviour (the motion of the attracted object). It does not matter if the second-order property is inert so long as we can maintain all three distinct features, for instance by admitting exercisings or manifestations as first-order properties—thus allowing first-order properties that are not picked out by what we class as occurrent—property terms.[7]

Like Mill, I recommend capacity talk wherever I find the analytic method in use. But unlike Mill, at least as Schmidt-Petri pictures him, I take this talk literally. The component features have capacities, the capacities are exercised, and the result of their joint operation is what happens. That is how the laws for the components—laws in my sense, ascribing capacities and describing their mode of operation, not laws in the regularity-among occurrent-properties sense—explain the result. What about Psillos? He tells us that the laws for the components "contribute to a full explanation" of what occurs when they operate together, also that these laws "govern" the complex effect without "covering" it. What then do these law claims say, and what sense is there to "govern" or even "explain" once both the covering-law story and the capacities story are rejected?

Alternatively Psillos suggests that complex laws could do the job. There would then have to be an open-ended collection of these laws, enough to cover every arrangement of contributing causes that ever occurs. The notion of regularities here is certainly strained; and if not regularities, what are the truth-makers for these laws? Besides that, I would still argue, as in *The Dappled World*, that even these need a *ceteris paribus* clause in front—"only so long as nothing interferes", where interference is a robust capacity concept.

NOTES

1. Or ascribes it to any object in the right circumstances instancing that property. Which way one puts it depends on how one wants to understand the metaphysics of capacities.
2. Note that my usage of these terms here differs from that of Psillos, who seems generally to use 'manifestation' to refer to what I call resultant behaviour.
3. As I note in *The Dappled World*, sometimes we do not have a nice summarizing word such as "attracts" for the manifestation or exercise of the capacity; hence the resort to "tries to X" where X is a canonical behaviour associated with the capacity.
4. Early Morrison paper
5. It was thus, as Psillos points out, a gross exaggeration on my part to say that the best evidence that one feature can cause another is that it does so, in the capacity sense of 'can'. This is good evidence only once we suppose (as in the Gravity Probe) that whatever the cause produces it does out of a stable capacity.
6. In this case it seems the manifestation occurs whenever the mass does. But that is not necessary—some capacities need triggering or manifest themselves only in special circumstances.
7. Nor do I object to the existence of "laws" that demand that different specific capacities occur together. On the other hand, I certainly would not admit them in order to explain why they occur together as I don't see why that needs explanation.