

The epistemic account of ceteris paribus conditions

Wolfgang Spohn

Abstract The paper focuses on interpreting ceteris paribus conditions as normal conditions. After discussing six basic problems for the explication of normal conditions and seven interpretations that do not well solve those problems I turn to what I call the epistemic account. According to it the normal is, roughly, the not unexpected. This is developed into a rigorous constructive account of normal conditions, which makes essential use of ranking theory and in particular allows to explain the phenomenon of multiply exceptional conditions. Finally, this static account is extended to a schematic dynamic model of how we may learn about those normal and (multiply) exceptional conditions.

Keywords Ceteris paribus laws · Normal conditions · Multiple exceptionality · Conditional belief · Ranking theory

1 Introduction

Apparently, scientists of many disciplines state laws qualified by a ceteris paribus (*cp*) clause; apparently, they intend to thereby make claims that can be true or false and somehow confirmed or disconfirmed; and apparently, philosophers of science have great difficulties in precisely saying what the contents or truth conditions of such claims are. Therefore, this paper will move away from truth conditions and attempt to motivate and elaborate an epistemic account of *cp* laws.

The motivating part comprises sections 2 – 4. First, reference to *cp* conditions is not a uniform affair; the label has many meanings. Here, I will restrict myself to interpreting *cp* conditions as normal conditions. Section 2 will briefly explain the content of this restriction. Section 3 will introduce six problems for specifying truth conditions for laws qualified by normal conditions; those six problems seem to lie at the heart of the relevant debate. Section 4 then identifies seven kinds of strategies of accounting for normal conditions and of dealing with those problems. All of the seven strategies seem to be at most partially satisfactory. Fully arguing this would require a

W. Spohn (✉)
Department of Philosophy, University of Konstanz, 78457 Konstanz, Germany
e-mail: wolfgang.spohn@uni-konstanz.de

long critical review. This is not my intention. Here, I will spell out my concerns with those strategies only in order to, hopefully sufficiently, motivate my search for an alternative strategy.

This strategy will consist in turning away from truth conditions and giving an epistemic account of the qualification by normal conditions, in which normal conditions will be, roughly, explicated as not unexpected conditions. I will explain why this strategic shift is able to improve the situation. Thus starts the constructive part of the paper. The theoretical vehicle for promoting the explication will not be (subjective) probability theory (which I will have put aside in section 4), but rather ranking theory, the basics of which are explained in section 5. Section 6 introduces the algebraic set-up that abstractly frames each discussion of *cp* laws. Sections 7 and 8 are the constructive core of the paper. The first precisely explicates epistemically interpreted normal and (multiply) exceptional conditions from the point of view of a given epistemic state or ranking function, while the second gives a simple schematic account of how we can reach such an epistemic state by learning from past observations. Section 9 will resume what has, and has not, been achieved here.

The basic ideas of this paper were first introduced in Spohn (2002), in quite obscure a way, as I retrospectively admit. I have elaborated on them in Spohn (2012, sect. 13.2), the relevant parts of which are here essentially reproduced in Sections 6–7 (with a few errors corrected). However, my presentation should be neater. Sections 2–4, the critical or motivating part should at the same time serve as a succinct overview of the theoretical options available in this field. In particular, by adding section 5 the paper is self-contained and thus gives direct access to my account of *cp* conditions, which is much less arduous than in Spohn (2012). Section 8 is new.

2 Four readings of *ceteris paribus* conditions

The problem of *cp* laws or conditions was introduced by Canfield and Lehrer (1961) into analytic philosophy of science. It took a while till the problem was taken seriously, Cartwright (1983), and Hempel (1988) being milestones. Nowadays it seems to be acknowledged as being one of the core problems of philosophy of science. The phenomenon of *cp* laws is very widespread (maybe it's even “*ceteris paribus* all the way down” to the fundamental laws of physics), and so is the uncertainty and inability to say what they are or claim, even after a very intensive debate in the last 10 or 20 years (see, e.g., Reutlinger et al. (2011), and Strößner (2013)).

Indeed, the label seems to encompass different kinds of phenomena, which should not be confused and may require different accounts. I even find different classifications. In my view, the best mnemonic is: *ceteris paribus* means *other things being equal* or *normal* or *absent* or *ideal*.

These are indeed four different understandings. “Other things being equal”, the literal translation of “*ceteris paribus*”, is a relational phrase. Equal to what? As long as the *relatum* is not specified, its meaning is incomplete. The intention certainly is that any other case may serve as standard of comparison. Then, however, the *cp* clause seems to amount to a general dictum like “equal causes, equal effects”, often called the general law of causality or the law of uniformity of nature. Discussing the status of these alleged laws is not our task here.

However, the usual reference point is not any given case and how it specifically is, but rather how things normally are. This reading of *ceteris paribus* conditions as normal conditions is an important one. Hooke's law is a pertinent example about well manufactured, homogeneous springs, handled with care and not exposed to too large forces, too high temperatures, and other exceptional conditions. When Carnap (1956, p. 69) stated that the reduction sentences associated with ordinary disposition predicates hold only *ceteris paribus* (thereby maneuvering the old discussion about dispositions to a dead end), he referred to normal conditions. And so on. What, though, *are* normal conditions? This will be our issue here. However, it gains further contour when we delineate it from further readings of *cp* conditions.

Equally common is the interpretation of “*ceteris paribus*” as the absence of disturbing factors. The law of falling bodies holds in the absence of further forces like air resistance, etc. Hooke's law is also a case in point; Canfield and Lehrer (1961)) presented it as posing a problem about disturbing factors. This shows that the classifications may well overlap. Still, the two interpretations must be distinguished; the absence of disturbances is not normal as such. Also, speaking of the absence of intervening variables presupposes that those variables have a natural value 0 that represents their absence. However, not all variables have such a value. Temperature, for instance, is always present; only 0K may be interpreted as the absence of temperature, and it is a most atypical condition. (Reutlinger et al. (2011, sect. 3.1) call this the exclusive sense of *cp* laws; however, in sect. 8 they go on to discuss normality theories as accounting for such exclusive *cp* laws. Thus, as mentioned, even classifications of the field do not fully agree.)

The fourth way interprets *ceteris paribus* as ideal conditions. This introduces another huge topic, idealization in the sciences. Again, the boundary to other interpretations is not sharp. The assumption that light moves with infinite velocity can only be seen as an idealization, whereas the absence of disturbing factors, such as the absence of air resistance, may, but need not be taken as an idealization. However, this interpretation introduces still further problems, which are not my concern here.

So, I will exclusively focus on normal conditions, and when I will use the ambiguous “*ceteris paribus*”, I will always refer to this reading as normal conditions, which is perhaps the most widely discussed one. It is confusing enough by itself; we should not further mix it up with the other readings of *cp* conditions just mentioned.

3 Six problems with explaining normal conditions

What is a law (statement)? At least, it seems, a true claim of the form “all *F*'s are *G*”. And something more, as the still on-going discussion about the mysterious modality of lawlikeness or nomicity shows. I am convinced that any progress on the notion of *cp* laws ultimately presupposes a firm conception of what laws are. With this connection firmly in mind, Lange (2002) and Woodward (2003, ch. 6), for instance, approach the problem of *cp* laws by amending lawlikeness through notions of stability or invariance; cf. Reutlinger et al. (2011, sect. 6). I have my own views; cf. Spohn (2012, ch. 12). However, I shall try to keep the discussion as independent from this background as possible; after all, our focus is “*cp*”, not “law”.

Maybe, we have already made the first, most natural step into the wrong direction. I said that a law at least makes a claim of the form “all F 's are G ”, which may be true or false. Hence, it is most natural to assume that a *cp* law “normally, F 's are G ” also makes a claim that may be true or false. After all, isn't science primarily the search for truth? So, almost the entire debate focused on the issue, on which I shall initially focus as well: what does a *cp* law really claim, what is its truth condition? Very rarely, the common presupposition of this debate was doubted. One instance is Earman and Roberts (1999) who despairingly concluded to put the notion of a *cp* law into the pragmatic waste-basket. I shall doubt it as well in the end.

However, in order to connect up with the existing discussion, let us share the common presupposition for the time being. My experience then is that there are basically six problems with explicating the truth conditions of *cp* laws, which somehow seem unsurmountable.

The first problem is that the claim made by “normally, F 's are G ” seems very vague, in fact *too indeterminate* to be useful, to definitely exclude any possibility whatsoever. Somehow, the explication of normal conditions should provide a more determinate meaning.

However, any proposal for more determinacy has the second problem that a *cp* claim is *too easily true*. If each or nearly each *cp* claim turns out to be true, this is not what we wanted. This would not be the way for science to give us substantial guidance.

The third problem is a special, more determinate version of the previous one; I call it the *opposition* problem. We shall see that various accounts admit that both, “*cp*, F 's are G ” and “*cp*, F 's are non- G ”, are true. This seems to be a knockdown argument against such accounts. Surely, under some (genetic) conditions men have blue eyes and under other conditions they have brown eyes. However, this doesn't entail that, *cp*, men have blue eyes and that, *cp*, men have brown eyes. We cannot have it both ways normally or *ceteris paribus*. The only *cp* law concerning eye color that comes to my mind is that, *cp*, men don't have red eyes.

Trying to avoid the Scylla of the second and the third problem takes us perilously close to the Charybdis of the fourth problem, namely that a *cp* claim is *too easily false*. Surely, a *cp* claim should be fallible; but if an explication tells us that we can never seriously maintain a *cp* claim, then this explication can't be right. (See also how Reutlinger et al. (2011, sect. 4) put this Scylla and Charybdis into the center of their survey.)

We will see various instantiations of these problems in the next section. However, we should first note how the Scylla and the Charybdis generate another deep dilemma. The fifth and perhaps biggest problem with *cp* laws is their *open texture*. The *cp* or normal conditions they refer to are usually taken from an open and entirely indeterminate space of possible conditions, a space of such-likes and don't-know-what's. Schiffer's (1991, p. 2) example “if a person wants something, then, all other things being equal, she'll take steps to get it” is a good psychological rule of thumb. However, countless conditions have to be satisfied in order for it to apply, and countless conditions may fail; we find no start and no end enumerating them.

We may try to escape this indeterminateness by referring to more determinate spaces of possible conditions. Then, however, we run either into the Scylla of the “too easily true” or into the Charybdis of the “too easily false”, as we will have occasion to observe. This in turn prompts the excuse that the more determinate space of possible

conditions was not meant so seriously, after all. So, welcome back to the “too indeterminate” problem.

The real problem here is this: Of course, we cannot discuss away the open texture of *cp* laws and the “too indeterminate” problem raised thereby. This is simply the epistemological condition we are in. Therefore, the hope can only be to say something reasonable and substantial about artificially determinate spaces of conditions, which then generalizes to indeterminate spaces in an illuminating way. However, if our account is plainly wrong for determinate spaces, this hope is thwarted; it remains unclear, then, how the account could improve by replacing the determinate by an indeterminate space of conditions. My main intent in this paper is to find a way to fulfill, not to thwart this hope.

Let me add a sixth problem, which is of a different kind; it will be important for my argument below. It seems obvious that there are not only normal and exceptional conditions, but also *doubly* or *multiply exceptional conditions*. There is not only Tweety, which is an exceptional bird, i.e., a penguin, which cannot fly. There is also Supertweety, which is an exceptional penguin which can fly and is hence not a normal, but rather a doubly exceptional bird. Each adequate account of normal conditions must allow for these distinctions. In the closely related literature on default logic or non-monotonic reasoning in general (see below), this point is well observed. The proper *cp* literature, however, does not seem concerned; for instance, the problem is not mentioned in the recent survey of Reutlinger et al. (2011). This seems to be a shortcoming.

4 Eight interpretations of normal conditions

Let us briefly look at the various accounts of *cp* or normal conditions and how they fare with respect to the problems just listed. The overview is useful as such, but it is also to motivate further elaboration of the account I favor. I am not sure whether my classification is optimal or complete; but at least the accounts I am familiar with seem to be fairly classified thereby.

There is first what I would like to call the *natural interpretation* according to which “*cp*, *F*'s are *G*” means “usually, normally, or mostly, *F*'s are *G*”. This is fine as far as it goes. However, it obviously has the “too indeterminate” problem. Scientists are not content with merely stating what usually happens; they want to be more precise and more specific. And philosophers are not happy with this explanation, which they can only take as the starting point of further analysis.

The natural interpretation naturally points to a statistical rendering, and the quantitatively tinged “mostly” does so more clearly than “usually” or “normally”. So, we secondly have the *statistical interpretation* according to which “*cp*, *F*'s are *G*” is rendered precise as “the statistical probability $P(G | F)$ of *F*'s being *G* is $> x \geq .5$ ”, or as some other probabilistic condition. Schurz (2001, 2002) defends and elaborates this interpretation (by providing an evolution-theoretic underpinning).

Still, I observe that it is widely taken as insufficient. What does statistical probability mean here? Could this be the limit of the relative frequency of *G*'s among *F*'s? Among the actual *F*'s? Hardly; there won't be infinitely many *F*'s. Among the counterfactual *F*'s? But how to constrain counterfactuality here? So, maybe the statistical probability rather is the propensity of *F*'s being *G*? No, *F*'s rather have a varying propensity of

being G under varying conditions. So, does the statistical probability of F 's being G , or the probability of a random F being G , average those propensities across the varying conditions? In order to assess normality and exceptionality, we then should somehow know the statistics for the occurrence of those conditions. But does such a statistics make sense? For instance, what is the statistics for the conditions under which springs behave according to Hooke's law? Moreover, the usual F 's are not the most F 's in the universe in any statistical sense; they rather are something like the most F 's we encounter. "Usual" usually means "usual for us" and is clearly an indexical; it is thus badly suited for statistics. I do not want to claim that these questions are unanswerable. However, they suggest that statistical probabilities are not available for many putative cp laws.

Therefore, a strong opposition against the statistical interpretation derives from the fact that there are many disciplines that state cp laws, but never understood themselves as statistical disciplines. Is Hooke's law at bottom a statistical law? This sounds odd. cp laws by far transcend the statistical sciences. Hempel (1988, pp. 152f.) paradigmatically wrote (although his provisoes should not be equated with cp clauses): "... a probabilistic construal of provisoes faces the difficulty that scientific theories do not, in general, provide probabilistic laws that would obviate the need for provisoes." This is not to reject the statistical interpretation; where it works, it's fine. But at least it seems too restricted; we have to account also for the cases where the statistical interpretation does not work. The main bulk of the literature attempts to do just this. For this reason I shall not further discuss the statistical interpretation.

There is, thirdly, what I call the *trivial interpretation* according to which " cp, F 's are G " means " F 's are G , unless they aren't". Of course, this is a caricature of an account, and it clearly falls prey to the "too easily true" and the opposition problem; in fact, it turns cp laws into analytic truths, for any F and G . Everyone rejects this interpretation. I only list it because everyone is in danger to slip into it; it is our no go area.

The first response to the challenge posed by Canfield and Lehrer (1961) was the fourth, the *eliminativistic interpretation*, as I call it; it is roughly to be found, e.g., in Stegmüller (1969, pp. 149ff.). According to it, " cp, F 's are G " means "under conditions C , all F 's are G ". So, " cp " is simply a lazy, but principally eliminable abbreviation for a usually very complex clause C explicitly stating conditions under which all F 's are G ; and cp claims simply are restricted universal claims. This won't do, though. This interpretation faces the "too easily false" problem; for any specific proposal for C actually envisaged the cp law is most likely false and not seriously maintained by us. In particular, it won't do as an analysis. Any given condition C is nothing but a doubtful hypothesis about what " cp " amounts to and can never be synonymous with " cp ". We may be lucky to hit upon the right C ; but we can hardly be confident.

Surely, though, we are so optimistic to think that *some* such C must hold good, even if we don't trust in any specification with which we could presently come up. and so the most common move is from the eliminativistic to the fifth, *existential interpretation*, as I call it; Reutlinger et al. (2011, sect. 5) call it the method of completers. According to it, " cp, F 's are G " amounts to the claim "there are suitable conditions C under which all F 's are G ". I find this idea elaborated, e.g., in Fodor (1991) and Pietroski and Rey (1995) and critically discussed in Schiffer (1991, sect. II), and Earman and Roberts (1999). The proponents seem to be aware of the pitfalls; therefore they spend large

efforts on restricting the existential quantification to *suitable* conditions C (which we need not scrutinize now).

As illuminating as those efforts are, they can't succeed in avoiding the "too easily true" problem, as noted already by Mott (1992). In particular, there will usually be suitable conditions C under which F 's are G and suitable conditions C' under which F 's are non- G . For instance, there are suitable conditions under which human eyes are blue and other suitable conditions under which human eyes are brown. So, the existential interpretation clearly faces the opposition problem. (See also the critical discussion in Reutlinger et al. (2011, sect. 5).)

One could have been warned. Existential quantifications over predicates, properties, or conditions always fail to do the job. This is a general lesson, which could be observed, for instance, already with the attempts of Kaila (1941) to save reduction sentences for disposition predicates from applying to too many objects (see also the description of the private discussion between Carnap and Kaila in Stegmüller (1970, pp. 221ff.)).

A sixth idea is to emphasize the explanatory role of *cp* laws. On one account, "*cp*, F 's are G " says something like "being F is a partial cause of, or plays a causal role for, being G ". This may be called the *causal interpretation*, which one finds endorsed, e.g., in Woodward (2002). On a similar account, "*cp*, F 's are G " roughly says that being F disposes to being G ; this may be called the *dispositional interpretation*. This view is attributable to Cartwright (1983) and others (cf. Reutlinger et al. (2011, sect. 7)). This idea at least indicates how *cp* laws can be used in scientific explanations. Woodward (2002), however, concludes that this idea rather speaks for neglecting *cp* clauses and focusing directly on the causal business. And again, both versions of this idea are plagued by the "too easily true" problem and even by the opposition problem. As with the existential interpretation, we will find conditions under which being F contributes or disposes to being G and other conditions under which being F contributes or disposes to being non- G . For instance, given daylight being yellow contributes or disposes to appearing yellow, but given appropriate artificial light being yellow contributes or disposes to appearing *not* yellow.

One might try to combine various interpretations. However, as far as I see, this does not help. We might add the observation that the statistical interpretation is perhaps able to deal with the problem of multiple exceptionality, and that the causal or dispositional interpretation need perhaps not be frightened by the open texture problem. But otherwise these two problems are not accounted for by the six proposals mentioned so far.

The preliminary conclusion is that the search for truth conditions of *cp* claims, to which the six interpretations hitherto discussed were devoted, was not successful. The sciences seem to continuously make claims implicitly or explicitly restricted by *cp* clauses, but we seem to be unable to say what these claims really are. This is disturbing; somehow our understanding of the matter is deeply wanting. Perhaps, the focus on truth conditions was already misguided.

A step away from this focus consists in noticing the similarity of normal conditions and default assumptions or of *cp* laws and default rules; generics, as treated in linguistics, play a similar role. This suggests exploiting the resources of default logic and non-monotonic reasoning for our topic – and results in what I call the *default interpretation*. The relevant literature is very rich, and the offers are too ramified to

allow a detailed discussion here. Given this richness it is disappointing how little philosophers of science have taken up this connection; to my knowledge, only Silverberg (1996) and Schurz (2001, 2002) have pursued this path.

In principle, I am full of sympathy with this line of investigation. Let me only add a few cautionary remarks. The primary interest in that field is to come up with a logic. It is part of all those logics that they avoid the opposition problem and allow for multiple exceptionality. However, the logic as such does not tell us what a default rule means; it only tells how to reason with it. The meaning issue is deferred to the semantics grounding the logic at hand. And there we find again a variety of explanations. Some (such as Reiter (1980)) assume default rules in the semantics as well and thus give, in effect, no explanation at all. Some (such as Veltman (1996)) explain default rules relative to a set of normal worlds and thus perhaps offer at least a systematization. The usual paradigm is still truth-conditional semantics; however, in which sense truth conditions are thereby provided remains a good question. Some explain default rules determinately in epistemic terms, but the conception of epistemic states varies from sets of sentences (Veltman (1996)) to ranking functions (Goldszmidt and Pearl (1996); Kern-Isberner and Thimm (2012)).

So, even though this type of interpretation goes into the right direction, it does not provide very clear guidance. Therefore, I would like to pursue here a determinately *epistemic interpretation*, which explicates *cp* laws and normal conditions directly by modeling the underlying epistemic states, and not indirectly via default logic. This is the eighth and final interpretation introduced here. Within the narrower *cp* literature it is particularly Schiffer (1991) and Earman and Roberts (1999) who recommend focusing on the epistemic or pragmatic role of *cp* clauses. However, they do not get beyond illuminating case studies; their general statements are rough and vague and are not backed up by any general epistemological theory. It is such a backing that I attempt to provide here.

However, “epistemic interpretation” is at present just a label without clear content. What is it supposed to mean? What does it say is the meaning of *cp* laws, if not a truth condition? And what might be the source of its alleged ability to succeed where the search for truth conditions fails? The answer to these questions is fundamental for the explicatory strategy pursued in the rest of the paper.

First, within the epistemic interpretation the meaning of an utterance is not a truth condition, but the mental state expressed by that utterance; this is the general strategy of expressivism. And if the utterance is a claim or an assertion, the mental state expressed is obviously an epistemic state (where any part of an overall epistemic state is an epistemic state as well). Thus, an epistemic interpretation of *cp* claims tries to say which epistemic states they express.

The epistemic interpretation may thus appear to be a big step away from common sense. However, one should note that expressivism in general is a semantic strategy that is taken seriously by many. In particular, it has strong defenders when it comes to an epistemic interpretation of (indicative) conditionals (cf. Adams (1975), Gärdenfors (1981), Edgington (1995)). This is a related topic, and therefore I have applied the same theoretical tools to the topic of conditionals in Spohn (2013) as I will do here.

So far, I have stated in abstract how the epistemic interpretation and its account of meaning differ from the truth-conditional strategy. Yet, one may wonder what is thereby gained. Of course, the epistemic interpretation gets substance only through

an account of epistemic states. However, after I have put aside any probabilistic ventures, the only common epistemic states left are beliefs or sets of beliefs. And if this is what gets expressed, there can't be any gain, because the content of a belief expressed by an assertion is just its truth condition. The shift to the epistemic interpretation would be pointless.

Therefore, it will be crucial to my argument that *cp* claims express not only beliefs, but also conditional beliefs. This by itself is not surprising; *cp* claims trust in their hypothesis only given normal conditions. The point, though, is that conditional beliefs have in general no truth conditions (see Spohn (2012, sect. 15.3)). This point is also at the heart of my epistemic account of conditionals in Spohn (2013). And this is precisely why the epistemic interpretation can be more successful than the search for truth conditions. After having laid out the constructive details these abstract remarks will be more intelligible.

To be a bit more specific: According to the epistemic interpretation I shall propose, exceptional conditions roughly are unexpected and normal conditions not unexpected. Thus, "*cp, F*'s are *G*" now means "given not unexpected conditions, *F*'s are *G*", i.e., this is the conditional belief expressed by the *cp* claim. As stated, this is the point that cannot be captured in terms of truth conditions. It may be natural to speak of expectations here, and one may try to distinguish beliefs and expectations (as, e.g., Gillies (2004) does). I don't see any need for doing so and will identify the two notions here; both consist in taking something to be true.

Thus I have laid out my key notions. As agents and scientists we continuously have beliefs, expectations, conjectures, and we do so unconditionally as well as conditionally, and we may or may not believe or expect the conditions to obtain in turn. This is a fundamental part of our cognitive enterprise, and *cp* clauses derive from the fact that we continuously hedge our beliefs and expectations by various more or less determinate conditions.

So, the way ahead of us seems clear. If we want to tackle these phenomena systematically, we must use systematic theories of (conditional) belief. It is my long-standing conviction that ranking theory is the most adequate account of conditional belief. I have amply defended this claim in Spohn (2012). Here I will simply proceed from it. So, let me first introduce the basics of ranking theory as far as needed here and then apply it to the special case of a hypothesis believed only under various conditions. Then we will also see how the epistemic interpretation fares with respect of the six problems in section 2, at which at least the first six interpretations mentioned above foundered in some way or other.

5 Basics of ranking theory

Probability theory, subjectively interpreted, is a theory of degrees of belief. However, as is well known, it is unable to represent belief; neither probability 1 nor probability $\geq 1 - \epsilon$ may be equated with belief, as the lottery paradox shows (and there is no good way getting around it). This is why I put aside subjective probabilities as well. Ranking theory does both. Its first concern is to represent belief. A static theory could do without degrees (see Hintikka (1962)). But if we want to account for the dynamics of belief, as ranking theory does, we have to reintroduce degrees of belief, which differ from

probabilities, even though they behave similarly in many respects. Let's see how the theory works (for all details see Spohn (2012, ch. 5)):

Just as in probability theory, we start with a space W of possibilities or possible worlds and some algebra \mathcal{A} of propositions over W . Then κ is a *negative ranking function* for \mathcal{A} iff κ is a function from \mathcal{A} into $\mathbb{N} \cup \{\infty\}$ (\mathbb{N} is the set of non-negative integers) such that for all $A, B \in \mathcal{A}$:

$$\kappa(W) = 0 \text{ and } \kappa(\emptyset) = \infty, \quad (1)$$

$$\kappa(A \cup B) = \min\{\kappa(A), \kappa(B)\} \quad (\text{the law of disjunction}). \quad (2)$$

A negative ranking function is to be interpreted as a *grading of disbelief* (this is why I call it negative despite its non-negative values). $\kappa(A) > 0$ says that A is *disbelieved* or *taken to be false*; and the larger $\kappa(A)$, the stronger the disbelief in A . $\kappa(A) = 0$ says that A is not disbelieved. So, (1) says that the tautology W is not disbelieved at all and that the contradiction \emptyset is maximally disbelieved. It may well be that both $\kappa(A) = 0$ and $\kappa(\bar{A}) = 0$; this represents neutrality or suspense of judgment w.r.t. A . Finally, belief in A is the same as disbelief in \bar{A} ; so, $\kappa(\bar{A}) > 0$ says that A is *believed* or *taken to be true*. This is how ranking theory represents belief.

(2) is the characteristic axiom saying that a disjunction cannot be less disbelieved than its disjuncts. (2) also says that the disjunction is not more strongly disbelieved than either disjunct. This is required, because the disjunction is logically weaker than the disjuncts. In this way, (2) embodies the deductive closure of belief (which means in terms of disbelief that, if you disbelieve any propositions, you also disbelieve any logically stronger propositions). In this interpretation (1) entails that there is a proposition, namely \emptyset , which is not believed, and thus that deductively closed belief is consistent.

This may seem to be a roundabout way of representing belief, and it has certainly hampered the reception of ranking theory. Of course, we may also define a *positive ranking function* β by

$$\beta(A) = \kappa(\bar{A}) \quad (3)$$

expressing positive belief; or we could axiomatize it directly. However, conditional belief, the dynamics of belief, indeed the entire mathematics runs more smoothly in terms of negative ranking functions. Moreover, the far-reaching analogy of probability theory stands out only in this way. So, I will continue in those terms and usually neglect the adjective “negative”.

(1) and (2) immediately imply:

$$\kappa(A) = 0 \text{ or } \kappa(\bar{A}) = 0 \quad (\text{the law of negation}). \quad (4)$$

(4) says that your beliefs must be consistent; you can't disbelieve (or believe) both A and \bar{A} .

If we are particularly dealing with infinite algebras, then we should also postulate an infinite version of axiom (2). However, there is no need to go into these formal details.

The all-important notion is the notion of conditional ranks. The *conditional (negative) rank* $\kappa(B | A)$ of B given A is defined as

$$\kappa(B | A) = \kappa(A \cap B) - \kappa(A), \quad (5)$$

provided $\kappa(A) < \infty$. This is tantamount to the quite intuitive *law of conjunction* that your degree of disbelief in a conjunction $A \cap B$ adds your distrust in B given A to your distrust in A .

So, as before $\kappa(B | A) > 0$ expresses disbelief in B given A ; $\kappa(B | A) = 0$ expresses that B might be true given A according to κ ; and $\kappa(\overline{B} | A) > 0$ expresses *belief in B given A* . We will continuously use this later on. Given definition (5), the fundamental law of disjunction (2) is equivalent to:

$$\kappa(B | A) = 0 \text{ or } \kappa(\overline{B} | A) = 0, \quad (6)$$

which is a conditional version of the law of negation (4) and thus requires that your conditional beliefs be consistent as well. Hence, the deductive closure of beliefs, which is entailed by (2), reduces to conditional consistency. In other words, ranking theory embodies nothing but the definition (5) of conditional ranks and conditional consistency (6) – and thus has very strong normative foundations.

The similarity of (5) with conditional probability is salient. It is therefore no surprise that the theorems for conditional ranks closely resemble those for conditional probability. There is, for instance, a ranking-theoretic version of Bayes' theorem, which seems to provide a basic form of the widely used inference to the best explanation, and which may even be more adequate than its probabilistic original, for the simple reason that this inference pattern, just as a *cp* law, is usually not perceived as having a probabilistic nature. (Cf. Spohn (2012, sect. 14.15) for details.)

The only theorem we shall need below is the following:

$$\text{if } \kappa(C | A) \leq \kappa(C | B), \text{ then } \kappa(C | A) \leq \kappa(C | A \cup B) \leq \kappa(C | B) \quad (7)$$

(law of disjunctive conditions)

This says that the conditional rank of a proposition C given a disjunctive condition lies between the conditional ranks of C given the disjuncts. It is easily proved (cf. Spohn (2012, theorem 5.23)), and familiar from probability theory.

The definition of conditional ranks is the crucial advance of ranking theory over its predecessors such as Shackle's (1969) functions of potential surprise or Cohen's (1977) Baconian probability. With their help we may also state adequate conditionalization or learning rules and thus a general dynamics of belief. However, for the things to come there is no need to develop ranking theory any further.

6 The algebraic set-up

So far, ranking functions were defined on some propositional algebra. For our present purposes it is useful to introduce a little bit more algebraic structure. It will consist in what I will call the target variable Y and the background variable Z .

What is a variable X in our context? Just what is called a random variable in probability theory, i.e., a measurable function from the space W of possibilities into some range W_X of possible values. Thus, a variable X represents some (complex) feature or magnitude which takes varying shapes (values in W_X) in different possibilities or worlds (in W).

Now, a *cp* claim first states a certain hypothesis that aims at, or is about, the target variable Y . The target variable Y may be compound in turn. For instance, Y may at once represent the force exerted on a spring and its extension, or it may describe demand, supply, and price of a certain good on a given market. The hypothesis then makes a claim about Y to the effect that only a subset of possible values in the range W_Y of Y will be realized, for instance, a subset that describes a functional relation between force and extension or between demand, supply, and price. We may further simplify matters and assume that the target variable Y takes just two values y_1 (= true) and y_0 (= false), so that the proposition $\{w \mid Y(w)=y_1\}$, or $\{Y=y_1\}$ for short, says that the hypothesis is true, i.e. *is* the hypothesis, and $\{Y=y_0\}$ is the negation of the hypothesis.

However, we do not simply believe in the hypothesis $\{Y=y_1\}$ about Y . There is a wider background of conditions which may be relevant to $\{Y=y_1\}$ and which we summarize in the background variable Z . Thus, Z may be very complex, consisting of many individual factors, which need not be further specified. Ultimately, Z is to represent the open-textured space of conditions, all those such-likes and don't-know-what's, which make our topic so obscure. However, we cannot formally reason with such an open-textured space. Therefore, I shall assume that the background variable Z is fixed and determinate; we will see later on how we can move from such fixed Z 's to open-textured backgrounds.

Possibilities thus reduce to pairs $\langle y, z \rangle$ of values of Y and Z , and propositions are sets of such pairs. There are what I call Y -propositions only about the variable Y that do not exclude anything about Z . The hypothesis $\{Y=y_1\}$ and its negation are Y -propositions, in fact the only substantial ones, after we have reduced the potentially very rich target variable Y to a binary one. And there are Z -propositions that do not exclude anything about Y ; we might also call them *background propositions*. We will also have to refer to most specific or logically strongest background propositions of the form $\{Z=z\}$, which I shall call *background states*.

Of course, there are also many mixed propositions about the relation between Y and Z . And our beliefs or epistemic attitudes concern the entire set-up. So far, the only point is to be able to distinguish the conceptual role of target and background variable. It should be clear, though, that this set-up is most general and covers any kind of *cp* law we might wish to formulate.

I have been ambiguous so far by sometimes referring to *cp* laws and sometimes to *cp* conditions; and the formal set-up does not disambiguate by itself. It is crucial to clear up this ambiguity. Talking about *cp* laws suggests that the hypothesis $\{Y=y_1\}$ at hand is a generalization, something like "all F 's are G ", which is hedged in some way. However, it is not the entire generalization that is embedded in *cp* or normal conditions; the idea is not that only in normal worlds all springs extend according to Hooke's law, without exception. Rather, the idea is that each single case is, or is not, embedded in normal conditions; a given spring extends lawfully, if it is situated in normal conditions. Hence, we must first study the single case, its behavior, and our beliefs about it. And then we may be able to generalize.

This is a general lesson, by the way. For instance, many attempted to say what causal laws are, which then instantiate in single cases and in particular causal relations. The more successful strategy, I believe, is to first clear up the causal relations in the single case, which may then generalize in a lawful way. (Cf. the endorsement of this view in Cartwright (1989, pp. 2ff.)) Again, if we want to understand the distinction between lawlike and accidental generalization, we also have to attend to the shape of our expectations about the single cases; or so I argued in Spohn (2012, ch. 12).

Compare this with the probabilistic case. There, too, at least if we move within a propensity interpretation, we do not start with a statistical law or generalization. Rather, we start with a statistical hypothesis about a single case, e.g., a single throw of a given die. And then we turn the hypothesis into a statistical law by generalizing to the infinite independent and identically distributed repetition of the single case, which generates a general distribution for infinitely many cases. Thus, the statistical law is derivative on the statistical hypothesis about the single case.

This is why I will understand the target variable Y , the hypothesis $\{Y=y_1\}$, and the background variable Z as referring to a *given single case*, not, e.g., to springs and markets in general, but to this spring in its present circumstances and to the present market situation for the relevant good. And we must study our beliefs or expectations about this single case. This is what I shall do in the next section.

Only in a second step we should consider how we reason about the generalization, i.e., about the (infinite) repetition of the set-up $\langle Y, Z \rangle$, i.e., about the (infinite) sequence of variables $Y_1, Z_1, Y_2, Z_2, \dots$, where each Y_i and Z_i is, respectively, a replica of Y and Z . This is just as with a stochastic process X_1, X_2, \dots consisting of an (infinite) repetition of a given random variable X .

In section 8 we will deal with the repetition $Y_1, Z_1, Y_2, Z_2, \dots$. First, however, we must attend to the single case as represented by Y and Z and to what normality might mean with respect to it. This single case may be anyone, say, the next one, that is not yet observed.

7 Normal and exceptional conditions: an explication

It lies in the nature of the epistemic interpretation that we do not consider the truth, but our beliefs about this set-up. So, let our epistemic state concerning this set-up be represented by the negative ranking function κ . What is the hypothesis maintained in κ ? It is, what we believe about the target variable Y . So:

$$\kappa(\{Y = y_0\}) > 0 \text{ (which says that } \{Y = y_1\} \text{ is believed).} \quad (8)$$

(8) entails $\kappa(\{Y=y_1\})=0$. This already describes the beliefs about Y as contained in κ (after we have reduced Y to a binary variable). But, of course, κ also holds beliefs about the background variable Z and sees Y and Z as connected; the beliefs about Y are conditional on Z in various ways. This is obviously the issue we have to scrutinize.

The basic observation, already sketched in Spohn (2002, sect. 4), is this: Even if $\{Y=y_1\}$ is unconditionally believed in κ according to (8), it need not be believed under all conditions. There will be some background condition or Z -proposition N that is unsurprising and not excluded, i.e., $\kappa(N)=0$. In this case we still have $\kappa(\{Y=y_0\} | N) > 0$; the

hypothesis $\{Y=y_1\}$ is still held under such a condition N . We may therefore call it a *normal* condition relative to the hypothesis at hand. However, there might also be a condition or Z -proposition E , given which $\{Y=y_1\}$ is no longer held, i.e., $\kappa(\{Y=y_0\} | E) = 0$. This entails $\kappa(E) = \kappa(\{Y=y_0\} \cap E) \geq \kappa(\{Y=y_0\})$; that is, such a background E is at least as disbelieved as the violation of the hypothesis. We may therefore call E an *exceptional* condition (relative to the given hypothesis), which is unexpected according to κ .

This is, I find, quite an appropriate schematic description of what actually goes on. We entertain a hypothesis, and this includes the expectation of conditions under which the hypothesis holds, or rather the denial of conditions under which it does not hold. When we nevertheless encounter a violation of the hypothesis, we are surprised, we start investigating how this was possible, and we find that some unexpected condition was realized under which we did not assume the hypothesis to hold, anyway. Or still more realistically, it is rather this investigation that elicits our tacit assumptions and expectations, the ranking function κ we are operating with. In this way, each ranking function representing beliefs in a hypothesis and about the relevant background automatically carries an aura of normal and exceptional conditions, which does not surface at the level of belief, but becomes visible only by looking at various levels of disappointed expectations, i.e., negative ranks and the conditional beliefs contained therein. This is the richer picture invoked by the epistemic interpretation of *cp* conditions.

However, we cannot stay with our preliminary explanations. We will have to go through a few insufficient explicative attempts, but only in order to acquire a better understanding of the logic, as it were, of normal and exceptional conditions (which are always to be understood relative to the given hypothesis) and thus of what I take in the end to be the right explanation. So far, I have given the following explication:

$$\text{the condition } N \text{ for } Z \text{ is } \textit{normal}^{nu} \text{ iff } \kappa(N) = 0, \text{ and} \quad (9)$$

$$\text{the condition } E \text{ for } Z \text{ is } \textit{exceptional}^* \text{ iff } \kappa(\{Y = y_0\} | E) = 0. \quad (10)$$

The superscript *nu* stands for “not unexpected” and indicates that we will distinguish further senses. (9) entails $\kappa(\{Y=y_0\} | N) > 0$, and (10) entails $\kappa(E) \geq \kappa(\{Y=y_0\})$.

(9) and (10) can’t be right, though. One should think that E and \bar{E} cannot both be exceptional. This is satisfied by “exceptional*”. If N is normal, should \bar{N} be not normal? This sounds implausible; for instance, concerning eye color being male is just as much a normal condition as being female. In any case, (9) allows both N and \bar{N} to be normal^{nu}. So far, so good. One should also think, though, that E is exceptional if and only if E is not normal. However, according to (9) and (10) non-normality^{nu} does not entail exceptionality*.

Should we therefore weaken exceptionality* and define E to be exceptional if $\kappa(E) > 0$? No, If still $\kappa(\{Y=y_0\} | E) > 0$, then, it seems, E is only unexpected, but not exceptional. So, we better weaken normality and define:

$$\text{the condition } N \text{ for } Z \text{ is } \textit{normal} \text{ iff } \kappa(\{Y = y_0\} | N) > 0. \quad (11)$$

However, (11) is as close to the trivial reading as one can get within our epistemic setting. It *defines* normal conditions to be such that given them the hypothesis is believed (though not necessarily true).

There are more problems. One problem that will finally set us on the right track is this: So far, there are many normal* and many exceptional* conditions. Intuitively, though, we might have thought that they are unique, that there is *the* normal and *the* exceptional condition. If this is not simply a careless use of the definite article, what could it mean? It might mean the strongest normal and the strongest exceptional condition. However, according to (11), if N and N' are normal*, $N \cap N'$ need not be. Hence, there is no strongest normal* condition. Similarly and contrary to expectations, if E and E' are both exceptional* according to (10), $E \cap E'$ need not be; we might have $\kappa(\{Y=y_0\} | E \cap E') > 0$. In this case, $E \cap E'$ would be normal*, although, intuitively, $E \cap E'$ looks *doubly* exceptional, because the judgment about $\{Y=y_1\}$ is reversed twice. So, (10) and (11) can't be quite right, either.

Perhaps, the definite article refers to the weakest, instead of the strongest, normal or, respectively, exceptional condition. This makes at least formal sense, since (10) and (11) entail:

$$\text{if } N \text{ and } N' \text{ are normal* conditions, so is } N \cup N', \quad (12)$$

$$\text{if } E \text{ and } E' \text{ are exceptional* conditions, so is } E \cup E'. \quad (13)$$

(12) and (13) are direct consequences of the law of disjunctive conditions (7). How does this help in finding the weakest normal or, respectively, exceptional condition? We know already that, taken literally, the tautology W is the weakest normal* condition. But this is not what could be meant.

We get a grip on the issue when we first focus on background *states*, i.e., on propositions of the form $\{Z=z\}$ saying that the background takes the maximally specific shape z . Those states are either normal* or exceptional*, and (12) and (13) allow to comprise them to a normal* and exceptional* condition:

$$\begin{aligned} N^* &= \cup \{ \{Z=z\} \mid \{Z=z\} \text{ is a normal* background state} \}, \\ E^* &= \cup \{ \{Z=z\} \mid \{Z=z\} \text{ is an exceptional* background state} \}. \end{aligned} \quad (14)$$

Thus, N^* and E^* are the weakest conditions of their kind that are pure in the sense of containing only background states of the same kind. And they are complementary: $E^* = \overline{N^*}$. So, do N^* and E^* deserve to be called *the* normal and *the* exceptional condition?

The answer is still no. N^* again amounts to the trivial reading in epistemic disguise. And we still have the problem that there are normal* conditions that are intuitively doubly or multiply exceptional. However, the observations (12) and (13) help us to the proposal I finally endorse (and that I had already intended in Spohn (2002, sect. 4)):

So far, we have found no reason to doubt the point that conditions under which the hypothesis $\{Y=y_1\}$ cannot be maintained are exceptional. Now, (13) entails that there is a weakest condition $E^{\geq 1} = \cup \{E \mid E \text{ is exceptional*}\}$, which is exceptional*, too (the superscript will soon explain itself). This is indeed a distinguished condition, which we may well call *the exceptional condition*. The distinction carries over to $N^0 = \overline{E^{\geq 1}}$, which may thus be called *the normal*

condition, as will be further justified below. N^0 is normal* and normal^{uu}; indeed N^0 is believed in κ , since $E^{\geq 1}$ is disbelieved.

More interesting is the observation that $E^* \subseteq E^{\geq 1}$ and $N^0 \subseteq N^*$. The reason is that N^0 will turn out to be that part of N^* not containing any fake normal, but actually multiply exceptional background states, as I am about to explain. This also means that N^0 cannot be denounced as a formalization of the trivial reading in epistemic disguise; it is essentially stronger.

Copying the argument for $E^{\geq 1}$, we are now able to account for the phenomenon of double exceptionality within the range of exceptionality. For $E \subseteq E^{\geq 1}$ we may call E *doubly exceptional* iff $\kappa(\{Y = y_0\} | E) > 0$, i.e., if the hypothesis $\{Y = y_1\}$ is in turn supported by E . According to (12), the disjunction of doubly exceptional conditions is doubly exceptional, too, so that we may define $E^{\geq 2} = \cup \{E \subseteq E^{\geq 1} | E \text{ is doubly exceptional}\}$ as *the doubly exceptional condition* (with respect to the hypothesis). Now the superscripts start to make sense.

Of course, we must not stop here; if there is double exceptionality, there can be triple exceptionality, and so on. How many degrees of exceptionality there are depends on the richness of the background variable Z and on the sophistication of our judgment about the hypothesis $\{Y = y_1\}$ (as represented by κ). In practice, there will be few degrees; it is, however, not so difficult to construct contrived examples in which degrees of exceptionality build up indefinitely.

In any case, the general definition is now clear:

$$\begin{aligned} \text{Set } E^{\geq 0} &= W \text{ (= the tautology or the full possibility of space).} \\ \text{For odd } n \text{ we define } E^{\geq n} &= \cup \{E \subseteq E^{\geq n-1} | \kappa(\{Y = y_0\} | E) = 0\}; \\ \text{for even } n \text{ we define } E^{\geq n} &= \cup \{E \subseteq E^{\geq n-1} | \kappa(\{Y = y_0\} | E) > 0\}. \end{aligned} \quad (15)$$

We may go on to define:

$$E^n = E^{\geq n} - E^{\geq n+1} \quad (16)$$

Thus, N^0 as explained above is the same as E^0 . I stick to the letter N in order to emphasize that exceptionality of degree 0 is in fact normality. The sequence $E^0 = N^0, E^1, E^2, \dots$ forms a partition of W , and it provides a classification of background states: the state $\{Z=z\}$ is *exceptional to degree n* (and *normal* if $n=0$) iff $\{Z=z\} \subseteq E^n$. Thereby, we finally arrive at a classification of all background conditions:

$$\begin{aligned} \text{A } Z\text{-proposition or condition } E \text{ for } Z \text{ is } \textit{exceptional to degree } n \\ \text{with respect to the hypothesis } \{Y = y_1\} \text{ iff } n \text{ is the minimal degree} \\ \text{of exceptionality among the states } \{Z = z\} \subseteq E. \text{ If } n \text{ is even, we have} \\ \kappa(\{Y = y_0\} | E) > 0; \text{ if } n \text{ is odd, we have } \kappa(\{Y = y_0\} | E) = 0. \\ \text{A condition exceptional to degree 0 is } \textit{normal}. \end{aligned} \quad (17)$$

The claims contained in (17) are entailed by (12) and (13) and thus ultimately by the law of disjunctive conditions (7). According to (17), *the* normal condition N^0 is believed in any case, while *a* normal condition N may even be

disbelieved and not normal^m; it must only be not as disbelieved as $E^{\geq 1}$. In fact, we have:

$$\{Z = z\} \subseteq N^0 \text{ if and only if } \kappa(\{Z = z\}) < \kappa(E^{\geq 1}). \quad (18)$$

(This corrects claim (13.9) in Spohn (2012, p. 316).)

(17) explains the mistake in our attempt (14) to define the normal and the exceptional condition as N^* and E^* . We have $N^* = \cup \{E^n \mid n \text{ is even}\}$ and $E^* = \cup \{E^n \mid n \text{ is odd}\}$. Hence, this attempt misclassified all conditions of an even degree of exceptionality, including the background states, as normal*.

Finally, degrees of exceptionality are nicely ordered by their negative ranks. We have:

$$\kappa(E^{\geq 0}) \leq \kappa(E^{\geq 1}) \leq \kappa(E^{\geq 2}) \leq \dots \text{ and } \kappa(E^n) = \kappa(E^{\geq n}) \text{ for all } n. \quad (19)$$

Indeed, the weak inequalities can almost always be replaced by strict inequalities; for details see Spohn (2012, pp. 315f.). This agrees with our intuition that the higher the degree of exceptionality of a condition, the more unexpected it is.

Let me stop here with what I take to be the logic of normal and exceptional conditions. Let me once more emphasize that I have no general hope of probabilistically approximating (12) – (19). By contrast, this chain of definitions and observations could be duplicated, as far as I see, within AGM belief revision theory on the basis of entrenchment orderings; this is so because the account above only refers to the ordinal features of ranking functions. However, already the extension considered in the next section works only within the cardinal framework of ranking functions.

8 How we may learn about normal and exceptional conditions

In the previous section I explained normal and exceptional conditions with respect to a given hypothesis and relative to a given epistemic state or ranking function κ . This state held beliefs or expectations about the next, not yet observed relevant single case. However, we may assume, just as we do in inductive logic, that our future expectations satisfy the postulate of symmetry. This entails that they are generalizable; they hold for any unobserved single case, indeed for any collection of n future cases, in the same way. In each case, we expect that normal conditions obtain and that the hypothesis holds given normal conditions; in the absence of special information about a certain case we have no reason to expect anything else.

The issue I want to pursue in this section is how we may acquire such an epistemic state distinguishing between normal and (multiply) exceptional conditions in the way explained. This means leaving the static perspective adopted so far and turning to a dynamic perspective, which considers a sequence of single cases and a belief formation process driven by the observation of an initial segment of that sequence.

So, we are now considering a more complex algebraic set-up, not one target variable Y and one background variable Z , characterizing a given single case, but rather a sequence Y_1, Y_2, \dots of target variables and a sequence Z_1, Z_2, \dots of background variables such that each $\langle Y_n, Z_n \rangle$ is a replica of the single case $\langle Y, Z \rangle$ dealt with so far. Let's suppose, as before, that Y is reduced to a binary variable and that Z ranges over

$\{z_1, \dots, z_m\}$, so that each Y_n and Z_n , respectively, takes the values $y_{n,0}, y_{n,1}$ and $z_{n,1}, \dots, z_{n,m}$.

We will then study a sequence of epistemic states or ranking functions $\kappa_0, \kappa_1, \kappa_2, \dots$, where κ_0 is the a priori state, which holds initial beliefs about the entire sequence $Y_1, Z_1, Y_2, Z_2, \dots$, and where κ_n ($n \geq 1$) is the state reached after observing $Y_1, Z_1, \dots, Y_n, Z_n$, which holds revised beliefs about the future $Y_{n+1}, Z_{n+1}, Y_{n+2}, Z_{n+2}, \dots$. Each κ_n is symmetric and distinguishes between normal and exceptional conditions in the way described in the previous section, not only for the next case Y_{n+1}, Z_{n+1} , but for each future case in the same way.

Our task now is to explain how this distinction may be learned, and for that purpose we have to describe the sequence $\kappa_0, \kappa_1, \kappa_2, \dots$ in more detail. I shall present a very simple and schematic learning model, which, however, seems realistic enough to be instructive. Of course, if there is one model, one may think about many others. I am content with proposing at least one.

First, we have to fix the initial state κ_0 . Let's assume:

$$\kappa_0(z_i) = 0 \quad \text{for } i = 1, \dots, m, \quad (20a)$$

$$\kappa_0(y_0|z_i) = a > 0 = \kappa_0(y_1|z_i) \quad \text{for } i = 1, \dots, m. \quad (20b)$$

The notational convention here is that z_i stands here for the value $z_{t,i}$ of any future background variable Z_t , i.e., for the proposition $\{Z_{t,i}=z_i\}$ ("future" is relative to the relevant epistemic state). This convention is admissible, since the assumptions about each future case are the same, due to symmetry. Likewise for y_1, y_0 . This helps reducing notation.

(20) completely defines the initial state. (20a) says that initially we are entirely ignorant and unopinionated concerning the background conditions. (20b) says that we initially trust in the hypothesis $\{Y=y_1\}$ with strength a ; let's call a the *initial credit*. These assumptions are unobjectionable; and they entail that initially there are no exceptional conditions.

Where does the initial credit, the initial conjecture of the hypothesis come from? κ_0 is only prior to the investigation at hand and may look back to rich experience. This may suggest many hypotheses for further scrutinization: where there is smoke, there is fire; the scarcer the supply, the higher the prices; and so on. This is one possible source of the initial credit.

I want to emphasize, though, that the initial credit often has a conceptual nature. Reduction sentences for dispositions – "whenever something is put in water, it dissolves iff it is soluble" – are a priori hypotheses on conceptual grounds. However, they are not unrevisably a priori or even analytic, as originally conceived. They are hedged by a *cp* clause, as Carnap (1956, p. 69) concluded, and thus only defeasibly a priori. This is also instructive from the point of view of concept formation. (For the notion of defeasible apriority cf. Spohn (2012, sect. 6.5 and 17.1).)

The crucial issue now is this: how do we learn and revise the initial κ_0 ? Of course, by observing the first n cases $Y_1, Z_1, \dots, Y_n, Z_n$ and by conditionalizing on the data observed. Let us call those data D_n . Hence, $\kappa_n = \kappa_0(\cdot | D_n)$. So, it all depends on those conditional ranks about which we have to make appropriate assumptions. Defining $\kappa_0(\cdot | D_n)$ directly is not advisable. It is much more perspicuous to explain one inductive

step after the other. Thus, assume that we have already arrived at $\kappa_n = \kappa_0(\cdot | D_n)$ and that we now observe $y_{n+1,h}, z_{n+1,i}$ to obtain in the $n+1$ st case for some $h = 0, 1$ and $i = 1, \dots, m$. Hence, $D_{n+1} = D_n \cap \{Y_{n+1} = y_{n+1,h}\} \cap \{Z_{n+1} = z_{n+1,i}\}$ and $\kappa_{n+1} = \kappa_0(\cdot | D_{n+1}) = \kappa_n(\cdot | \{Y_{n+1} = y_{n+1,h}\} \cap \{Z_{n+1} = z_{n+1,i}\})$.

What we have to fix, then, is only the step from κ_n to κ_{n+1} . I propose the following rules, which I shall explain afterwards:

If $y_{n+1,1}, z_{n+1,i}$ is observed in the $n+1$ st case, then :

$$\begin{aligned}
 & \text{(a) } \kappa_{n+1}(z_i) = \kappa_n(z_i) - \nu_1, \\
 & \text{(b) } \kappa_{n+1}(z_j) = \kappa_n(z_j) + b - \nu_1 \text{ for } j \neq i (b \geq 1), \\
 & \text{(c) } \kappa_{n+1}(y_1 | z_i) = \kappa_n(y_1 | z_i) - \nu_2, \\
 & \text{(d) } \kappa_{n+1}(y_0 | z_i) = \kappa_n(y_0 | z_i) + 1 - \nu_2, \\
 & \text{(e) } \kappa_{n+1}(y_h | z_j) = \kappa_n(y_h | z_j) \text{ for } h = 0, 1, j \neq i.
 \end{aligned} \tag{21}$$

If $y_{n+1,0}, z_{n+1,i}$ is observed in the $n+1$ st case, then

$$\begin{aligned}
 & \text{(a) } \kappa_{n+1}(z_i) = \kappa_n(z_i) - \nu_3, \\
 & \text{(b) } \kappa_{n+1}(z_j) = \kappa_n(z_j) + 1 - \nu_3 \text{ for } j \neq i, \\
 & \text{(c) } \kappa_{n+1}(y_1 | z_i) = \kappa_n(y_1 | z_i) + c - \nu_4 (c \geq 1), \\
 & \text{(d) } \kappa_{n+1}(y_0 | z_i) = \kappa_n(y_0 | z_i) - \nu_4, \\
 & \text{(e) } \kappa_{n+1}(y_h | z_j) = \kappa_n(y_h | z_j) \text{ for } h = 0, 1, j \neq i.
 \end{aligned} \tag{22}$$

There are various things to explain about (21) and (22). First, those ν_r ($r = 1, \dots, 4$) are normalizing factors without which κ_{n+1} would not be a negative ranking function according to (1) and (2). Thus, $\nu_1 = \min_{j \neq i} \{\kappa_n(z_i), \kappa_n(z_j) + b\}$, so that $\kappa_{n+1}(z_k) = 0$ for at least one $k = 1, \dots, m$. And $\nu_2 = \min \{\kappa_n(y_0 | z_i) + 1, \kappa_n(y_1 | z_i)\}$. And so on. One may neglect those normalizing factors; this is not where the music plays.

Now, look at (a) and (b) of (21) and (22). They say that z_i , which is observed in the $n+1$ st case, improves its credibility as compared with the z_j ($j \neq i$); that is, the degree of disbelief in the latter raises according to (b). So, the ranks simply follow the observed absolute frequencies; the more often some background state occurs, the more credible it is.

However, in (22b) I add 1, whereas in (21b) I add $b \geq 1$. This is just to add some flexibility to the model; one may think about various other flexibilities. This one, though, is not arbitrary. One might call b a *normality bonus*. Note that in the case of (21b) $z_{n+1,i}$ is a background instance supporting the hypothesis $\{Y = y_1\}$, whereas in the case of (22b) it is an instance defeating the hypothesis. So, the normality bonus b expresses that we more strongly reckon with supporting than with defeating instances. (Of course, if $b = 1$, there is no bonus.)

It's the other way around with (c) and (d) of (21) and (22). In both cases, they say that the observed value of the target variable Y increases its credibility given the observed background state; or rather, the not observed value of Y increases its incredibility according to (21d) and (22c). This increase is 1 in the supporting case (21d) and $c \geq 1$ in the defeating case (22c). Again, though, c is not an arbitrary parameter. One may call c an *exceptionality alert*, which expresses that the disconfirming force of a

surprising defeating background instance is stronger than the confirming force of an unsurprising supporting instance. (If $c = 1$, there is no alert.)

Finally, (21e) and (22e) say that our belief in the hypothesis given a not observed background state $z_j \neq z_i$ is not affected at all by the observation of z_i ; it simply stays the same.

This concludes my assumptions about the learning step from κ_n to κ_{n+1} . Basically, they assume that we learn both about the background states and about the relation between the background states and the hypothesis by raising the credibility of the observed background state and of the observed relation – where this raising may be governed by various parameters. This is certainly a correct description of what goes on in actual cases, even though actual cases need not proceed as mechanically as (21) and (22) suggest.

From (20) – (22) we can infer the shape of our posterior epistemic state κ_n after n observations. It only depends on the observed absolute frequencies. So, let be:

$$\begin{aligned}
 p_i &= \text{the number of supporting observations of the form } y_1, z_i (i = 1, \dots, m), \\
 q_i &= \text{the number of defeating observations of the form } y_0, z_i (i = 1, \dots, m), \\
 n_i &= p_i + q_i = \text{the number of occurrences of } z_i, \\
 p &= \sum p_i = \text{the number of occurrences of } y_1, \\
 q &= \sum q_i = \text{the number of occurrences of } y_0, \\
 n &= p + q = \text{the number of observed cases.}
 \end{aligned} \tag{23}$$

It is quite straightforward then to add up the cumulative effect of the n learning steps according to (21) and (22). We get:

$$\begin{aligned}
 \text{(a) } \kappa_n(z_i) &= b(p - p_i) + q - q_i - \nu \quad (i = 1, \dots, m), \\
 \text{(b) } \kappa_n(y_0|z_i) &= a + p_i - \nu' \quad (i = 1, \dots, m), \\
 \text{(c) } \kappa_n(y_1|z_i) &= cq_i - \nu' \quad (i = 1, \dots, m).
 \end{aligned} \tag{24}$$

Here, ν and ν' are again normalizing factors: $\nu = \min_i [b(p - p_i) + q - q_i]$, and $\nu' = \min [a + p_i, cq_i]$, depending on i . κ_n is completely specified by (24) for any future case.

So, the hypothesis $\{Y = y_1\}$ is believed in κ_n given any of the background states z_i for which $a + p_i > cq_i$, and not believed given z_j if $a + p_j \leq cq_j$. The latter background states z_j are clearly exceptional according to the explication (17). But are all of the former z_i normal according to (17)? This is not guaranteed, but it holds under natural conditions. More precisely, we have:

For all $i, j = 1, \dots, m$, if $n_i \geq n_j, q_j \geq q_i$, and at least one of the two inequalities is strict, then $\kappa_n(z_i) < \kappa_n(z_j)$. Hence, if for all $i, j = 1, \dots, m$, for which $\kappa_n(y_0|z_i) > 0 = \kappa_n(y_0|z_j)$, we have $n_i \geq n_j$ and $q_j \geq q_i$, where at least one of these inequalities is strict, then all of those z_i are normal and all of those z_j are exceptional to degree 1. (25)

Proof According to (24a) we have $\kappa_n(z_i) < \kappa_n(z_j)$ iff $b(p - p_i) + q - q_i - \nu < b(p - p_j) + q - q_j - \nu$ iff $q_j - q_i < b(p_i - p_j)$ iff $q_j - q_i < b(n_i - q_i - n_j + q_j)$ iff $(b - 1)(q_i - q_j) < b(n_i - n_j)$. The latter follows from the premises of the first claim of (25), since $b \geq 1$. From this the second claim of (25) follows immediately with the help of (18).

What is the upshot of (25)? According to (24b-c), κ_n maintains the hypothesis $\{Y = y_1\}$, given those z_i for which $a + p_i > cq_i$ and does not maintain the hypothesis given

those z_j for which $cq_j \geq a + p_j$. And then all those z_i are classified as normal, and all those z_j as exceptional to degree 1, by our explication (17), provided the absolute frequencies are well-behaved, i.e., all of the normal background states occur more often than the exceptional ones and show less violations of the hypothesis. This seems to be a weak and reasonable presupposition for being classified, respectively, as normal and as exceptional.

So far, no higher degrees of exceptionality were generated by our simple learning model. This is because we have considered only one undifferentiated background variable Z . Differentiating the background leads to a more sophisticated view of exceptionality, as I will now at least suggest.

In the picture above we started maintaining a hypothesis, we cautiously hedged it by a *cp* clause, and we learned to distinguish normal and exceptional conditions, but we looked only at one type Z of background. Of course, we want to do better and to arrive at more reliable hypotheses with less exceptions; violations were even observed under normal conditions, and the hypothesis was sometimes satisfied even given apparently exceptional conditions. Of course, we wonder then how this comes about, and the natural idea is that these variations depend on still further background factors. We come up with a plausible conjecture what those factors might be, and then we run our observations and experiments once more attending to the refined background.

The result will still not be fully satisfactory; exceptions will still persist. So, we do it all over again, with a still more refined background. And so on. This may go on indefinitely; only rarely will the open-textured background be exhausted by a determinate list of background variables.

I think this is a relatively fair schematic description of the actual inductive practice; it is only unfair in describing us as passive observers rather than active experimentators. Can we duplicate it in our simple learning model? Certainly. We could start considering two or more background variables and state more sophisticated learning rules instead of (21) and (22). However, things quickly get tedious and imperspicuous.

There is a simpler schematic route to multiple exceptionality: We may have run our observations and experiments regarding the target variable Y and the background Z long enough for building up strong trust in our classification of normal and exceptional background states. So, we come up with a new hypothesis: if $\{Z = z_i\}$ is normal, then $Y = y_1$, and if $\{Z = z_j\}$ is exceptional, then $Y = y_0$. And we define a new target variable Y' which takes value y'_1 , if that new hypothesis is satisfied, and value y'_0 , if not. Of course, we hedge this hypothesis as well and conjecture it to depend on a new background variable Z' .

And so, we run our learning algorithm (20) – (22) with the new variables Y' and Z' . Thereby, we will come to classify some new background states z'_k as normal and some other z'_l as exceptional. If z_i is normal in the old set-up and z'_k is normal in the new set-up, z_i, z'_k is normal in the combined set-up with respect to the original hypothesis $\{Y = y_1\}$. If z_j is exceptional in the old set-up and z'_l is exceptional in the new set-up, z_j, z'_l is doubly exceptional in the combined set-up. And if exactly one is exceptional in its respective set-up, then z_j, z'_l singly exceptional in the combined set-up.

This may suffice as indication of how the learning rules might extend onto more complex backgrounds and higher degrees of exceptionality. No doubt, this section was much too short and could and should be turned into a systematic investigation. I am content, though, with having given an example of how an appropriately working

learning mechanism might be devised within the ranking-theoretic framework. I don't see how something similar could be provided within any other framework.

9 What have we achieved?

This was my constructive explication of the epistemic account. Let us check whether it satisfies our expectations. First, how does it fare with respect to the four basic problems of section 2? The “too vague” problem is avoided; the account is as rigorous as ranking theory itself. The opposition problem is avoided, too. One can't simultaneously maintain that, *cp*, an *F* is *G* and that, *cp*, an *F* is not *G*, since the normal condition is believed according to (18).

The “too easily true or false” problem is, in a way, trivially avoided, since *cp* claims are so far not assigned truth conditions at all. Of course, this is not a good response. However, concerning the “too easily true” problem, I have indicated how the epistemic account avoids the pitfalls of the trivial and the existential interpretation; it does so by giving the *cp* clause the stronger reading of the normal condition N^0 and not the trivial reading N^* or the existential reading “there are (suitable) conditions under which the hypothesis is held”. And concerning the “too easily false” problem, I have not explicated making a *cp* claim as simply making a claim we cannot be serious about. Rather, I have explicated it as a claim which we do trust in each instance, and have well described its inductive embedding, i.e., the aura of normal and exceptional conditions and its potential changes.

Moreover, our way from N^* and N^0 brought along a solution of the problem of multiple exceptionality. I have not seen anything of this kind in the proper *cp* literature. It should be compared, though, with the attempts in the default logic literature at dealing with hierarchies of default rules and thus with multiple exceptionality, and in particular with the ranking-theoretic neighbors Goldszmidt and Pearl (1996) and Kern-Isberner and Thimm (2012).

However, the greatest advantage of the epistemic over the other accounts seems to me to be the following: Other accounts appeared to have some plausibility when they are referred to open-textured background conditions. However, when they are related to a fixed and determinate background such as provided by the variable *Z*, they all turn out to be untenable. Their initial plausibility thus turns out to be unfounded. This is different with the epistemic account. It says something substantial in the case of a determinate background. And my dynamic learning story indicated how the account deals with determinately extended backgrounds and may thus even grow into open-textured backgrounds.

The final crucial question is still with us. Since the first six interpretations were unable to provide adequate truth conditions for a *cp* claim, I proposed the epistemic interpretation, which gave up on truth conditions. According to it, a statement hedged by a *cp* clause does not make a truth-evaluable claim, but only expresses the speaker's beliefs or epistemic state. I have explained why this can be a successful move. However, there should be some way back to truth conditions; it would be incredible to deny them entirely. Is there such a way? Yes, the expression of epistemic states is at least partially truth-evaluable, and the extent to which it is may be constructively investigated. For instance, unconditional beliefs are clearly truth-evaluable; they are

just as true or false as the propositions that are their content. So, which beliefs does a *cp* claim (about a single case) express?

According to (17), it expresses the unconditional belief in the normal condition N^0 , since $\kappa(\overline{N^0}) > 0$. Moreover, it characteristically expresses the conditional belief $\kappa(\{Y = y_0\} | N^0) > 0$ in the hypothesis $\{Y = y_1\}$ given the normal condition N^0 . Both entail the belief in $\{Y = y_1\} \cap N^0$, which is thus expressed as well, and may be true or false.

The rub is that the conditional belief in $\{Y = y_1\}$ given N^0 itself is not truth-evaluable; it only entails, but is not equivalent with, the truth-evaluable belief in the material implication $N^0 \rightarrow \{Y = y_1\}$. Thus, the issue “Which factual claim is associated with a *cp* statement?” is a peculiarly mixed affair.

This is not the place for further elaborating this issue; let me just refer to the so-called objectivization theory for ranking theory, which is developed in Spohn (2012, ch. 15) and which explores the extent to which features of ranking functions are objectivizable or truth-evaluable. There is a proof for the above stated fact that conditional belief is not generally objectivizable (cf. Spohn (2012, pp. 449f.)). It is also shown, however, that causal beliefs, which form a special class of conditional beliefs, *are* objectivizable given some weak assumptions. So, positive results are in reach.

However, I admit that this objectivization theory does not specifically extend to the conditional beliefs expressed by *cp* claims and to the hierarchy of multiply exceptional conditions underneath. Such an extension looks feasible, but I cannot report any positive results. In any case, I am convinced that, if we want to associate truth conditions with *cp* claims, we have to take the arduous path of the epistemic account and its accompanying objectivization theory.

References

- Adams, E. W. (1975). *The Logic of Conditionals*. Dordrecht: Reidel.
- Canfield, J., & Lehrer, K. (1961). Note on Prediction and Induction. *Philosophy of Science*, 28, 204–211.
- Carnap, R. (1956). The Methodological Character of Theoretical Concepts. In H. Feigl & M. Scriven (Eds.), *Minnesota Studies in the Philosophy of Science* (Vol. I, pp. 38–76). Minneapolis: University of Minnesota Press.
- Cartwright, N. (1983). *How the Laws of Physics Lie*. Oxford: Clarendon.
- Cartwright, N. (1989). *Nature's Capacities and Their Measurement*. Oxford: Clarendon.
- Cohen, L. J. (1977). *The Probable and the Provable*. Oxford: Oxford University Press.
- Earman, J., & Roberts, J. (1999). *Ceteris Paribus*, There is No Problem of Provisos. *Synthese*, 118, 439–478.
- Edgington, D. (1995). On Conditionals. *Mind*, 104, 235–327.
- Fodor, J. A. (1991). You Can Fool Some of the People All the Time, Every Else Being Equal: Hedged Laws and Psychological Explanations. *Mind*, 100, 19–34.
- Gärdenfors, P. (1981). An Epistemic Approach to Conditionals. *American Philosophical Quarterly*, 18, 203–211.
- Gillies, A. S. (2004). New Foundations for Epistemic Change. *Synthese*, 138, 1–48.
- Goldszmidt, M., & Pearl, J. (1996). Qualitative Probabilities for Default Reasoning, Belief Revision, and Causal Modeling. *Artificial Intelligence*, 84, 57–112.
- Hempel, C. G. (1988). Provisos: A Problem Concerning the Inferential Function of Scientific Theories. *Erkenntnis*, 28, 147–164.
- Hintikka, J. (1962). *Knowledge and Belief*. New York: Cornell University Press.
- Kaila, E. (1941). Über den physikalischen Realitätsbegriff. *Acta Philosophica Fennica*, 3, 33–34.
- Kern-Isberner, G., & Thimm, M. (2012). A Ranking Semantics for First-Order Conditionals, in: L. De Raedt et al. (eds.) *ECAI 2012*, Open Access, IOS Press, pp. 456–461.

- Lange, M. (2002). Who's Afraid of *Ceteris Paribus* Laws? Or: How I Learned to Stop Worrying and Love Them. *Erkenntnis*, 57, 407–423.
- Mott, P. (1992). Fodor and *Ceteris Paribus* Laws. *Mind*, 101, 335–346.
- Pietroski, P., & Rey, G. (1995). When Other Things Aren't Equal: Saving *Ceteris Paribus* Laws from Vacuity. *British Journal for the Philosophy of Science*, 46, 81–110.
- Reiter, R. (1980). A Logic for Default Reasoning. *Artificial Intelligence*, 13, 81–132.
- Reutlinger, A., Schurz, G., Hüttemann, A. (2011). *Ceteris Paribus* Laws, *Stanford Encyclopedia of Philosophy*, <http://plato.stanford.edu/entries/ceteris-paribus/>
- Schiffer, S. (1991). *Ceteris Paribus* Laws. *Mind*, 100, 1–17.
- Schurz, G. (2001). What is 'Normal'? An Evolution-Theoretic Foundation of Normic Laws and Their Relation to Statistical Normality. *Philosophy of Science*, 28, 476–497.
- Schurz, G. (2002). *Ceteris Paribus* Laws: Classification and Deconstruction. *Erkenntnis*, 57, 351–372.
- Shackle, G. L.S. (1969). *Decision, Order, and Time in Human Affairs*, Cambridge: Cambridge University Press, 2nd ed.
- Silverberg, A. (1996). Psychological Laws and Non-Monotonic Logic. *Erkenntnis*, 44, 199–224.
- Spohn, W. (2002). Laws, *Ceteris Paribus* Conditions, and the Dynamics of Belief. *Erkenntnis* 57, 373–394; also in: J. Earman, C. Glymour, S. Mitchell (eds.), *Ceteris Paribus Laws*, Dordrecht: Kluwer, 2002, pp. 97–118
- Spohn, W. (2012). *The Laws of Belief: Ranking Theory and Its Philosophical Applications*. Oxford: Oxford University Press.
- Spohn, W. (2013). A Ranking-Theoretic Approach to Conditionals. *Cognitive Science*, 37, 1074–1106.
- Stegmüller, W. (1969). *Probleme und Resultate der Wissenschaftstheorie und Analytischen Philosophie, Band I, Wissenschaftliche Erklärung und Begründung*. Berlin: Springer.
- Stegmüller, W. (1970). *Probleme und Resultate der Wissenschaftstheorie und Analytischen Philosophie, Band II, Theorie und Erfahrung, I*. Berlin: Springer.
- Ströbner, C. (2013). *Die Logik der Normalität: Untersuchungen zur Semantik von Normalitätssurteilen*. Berlin: de Gruyter.
- Veltman, F. (1996). Defaults in Update Semantics. *Journal of Philosophical Logic*, 25, 221–261.
- Woodward, J. (2002). There Is No Such Thing as a *Ceteris Paribus* Law. *Erkenntnis*, 57, 303–328.
- Woodward, J. (2003). *Making Things Happen. A Theory of Causal Explanation*. Oxford: Oxford University Press.