16

True Exceptions: Defeasibility and Particularism

*Bruno Celano**

A. Introduction

Sometimes, kinds of cases which do in fact fall under the antecedent of a conditional norm are reckoned recalcitrant, i.e. although they fall under the antecedent of the norm we do not wish to allow the consequence to follow. In such cases, we sometimes say that we are abandoning, or discarding, the norm. We concede, that is, that the alleged norm was wrong (or, if you wish, that it was no norm at all). At other times, however, it is claimed that the norm is a defeasible one. Granted, the case at hand is one of those in which the norm is defeated; but this, it is implied, does not amount to a wholesale abandonment of the norm itself. Claims to the effect that a given norm is defeasible, I take it, are different from claims of the former kind, i.e. claims to the effect that an alleged norm is simply wrong (or that it is no norm at all). Being defeasible, the norm somehow survives the impact of such recalcitrant cases. Though somehow revised, amended, qualified, the norm, it is assumed, remains in place: it is still the same norm.

I shall call this 'the identity assumption'. The identity assumption is part and parcel of the notion that some norms are not barely wrong, but defeasible. (This should be taken as a stipulation.)[1] There are various ways of substantiating the identity assumption, corresponding to some main ways of moulding the concept of defeasibility (of norms).

I shall challenge the identity assumption. Or, rather, I shall distinguish two different versions of it, and argue that one of them is hollow. I shall challenge it by reviewing, and rejecting, one of the main ways in which it can be substantiated, namely, a specificationist strategy for dealing with norm conflicts and inappropriate normative

* Professor of Legal Philosophy, University of Palermo. An earlier version of this paper was presented at the workshop on Defeasibilism held at Oriel College, Oxford, on 10–11 March 2008. I am grateful to the participants to the workshop, and to Riccardo Guastini, Giorgio Maniaci, and Giorgio Pino, for comments and criticism.

[1] The term 'defeasibility' designates—in contemporary inquiries dealing with practical reasoning (in morals, law, etc.)—different phenomena (recent attempts at sorting them out may be found in, e.g. J.L. Rodríguez and G. Sucar, 'Las trampas de la derrotabilidad. Niveles de análisis de la indeterminación del derecho' in J.C. Bayón and J.L. Rodríguez, *Relevancia normativa en la justificación de las decisiones judiciales* (Bogotá: Universidad Externado de Colombia, 2003), 115–16; J.C. Bayón, 'Por qué es derrotable el razonamiento jurídico?' in J.C. Bayón and J.L. Rodríguez, *Relevancia normativa en la justificación de las decisiones judiciales*; J.L. Rodríguez, *Lógica de los sistemas jurídicos* (Madrid: Centro de estudios políticos y constitucionales, 2002), 356–99; J. Hage, 'Law and Defeasibility', *Artificial Intelligence and Law*, 11 (2003); E. Bulygin, 'Review of Jaap Hage's Law and Defeasibility', *Artificial Intelligence and Law*, 11 (2003); H. Prakken and G. Sartor, 'The Three Faces of Defeasibility in Law', *Ratio Juris*, 17 (2004). In this paper, my topic is defeasibility of conditional norms, and, thus, of the inferences in which they may occur as premises. This, by itself, doesn't say much, of course. The specific phenomenon that I am concerned with will be apparent from the text.

verdicts (sections B and C). This will lead us to taking a stand in the generalism vs particularism debate (section D). Rejection of the identity assumption leads, when conjoined to an awareness of the phenomena underlying defeasibility claims, to a version of particularism.

Does this particularist stance necessarily involve rejection of any normative generalization? The answer to this question, I shall argue, turns on whether some sense can be made of the notion of a kind of case being 'normal' (section F). Here, a second version of the identity assumption may perhaps be endorsed, thus reviving the defeasibilist project. At the heart of defeasibility claims, so understood, is the idea that norms may be shot through with true, as opposed to merely prima facie, exceptions (section E).

Two caveats are in place here. First, what my argument in this paper is about are norms, not norm formulations (i.e. sentences which are taken to be expressing norms), nor the ways in which we get, from norm formulations, to norms which are taken to be expressed by them. Thus, the issues I shall be dealing with are not issues of interpretation—namely, issues concerning the ascription of meaning to sentences. I grant that much of what is discussed under the heading 'defeasibility in law' may be understood, and properly recast, as a cluster of issues about the ways in which meaning (often, different meanings) may be ascribed to legal sentences.[2] Here, much depends on whether one has a sceptical view of legal interpretation, that is, one leaving no room at all (except in the realm of errors and myths) for the idea of an interpretation of a given norm formulation being the right (or, perhaps, just a wrong) one. If all thoughts of criteria of right and wrong in interpretation are abandoned, if, namely, it is claimed that norm formulations may be interpreted and reinterpreted *ad libitum* so as to draw, from them, opposite conclusions, then no room is left for arguments about whether norms should be taken to be defeasible, in what sense, etc. These issues simply become otiose.[3] In what follows, I am assuming that some such room is left open—that we may sensibly argue about the forms and ways of good and bad inferences involving norms. But I shall not argue for this assumption. Specifically, I do not say anything in defence of the implied claim that, to some extent and in some cases at least, it may be right (or wrong, as the case may be) to claim that norm N is the norm that a given norm formulation expresses. What I am concerned with is the rational reconstruction of (good and bad) forms of inferences involving norms. Thus, the argument is normative (I mean epistemic normativity) throughout.

Second, I shall discuss 'norms' generally, leaving deliberately aside the issue whether we are dealing with either moral or legal norms, or both (or maybe other kinds of norms as well). This is because what I am interested in are abstract relations between normative contents, and the ways it is possible to understand or handle them in the face of the relevant phenomena (i.e. norm conflicts, inappropriate verdicts). *Vis-à-vis* this concern, the panoply of theoretical assumptions required in order to frame our inquiry as one concerning moral rather than legal norms, or vice versa, or the relations between the two—assumptions concerning the nature of law or of morality, or how can law incorporate or refer to morality, or how can morality leave room for the law, and so

[2] R. Guastini, 'Defeasibility, Axiological Gaps, and Interpretation', in this volume, ch. 9. See also J.L. Rodríguez, 'La derrotabilidad de las normas jurídicas' in J.C. Bayón and J.L. Rodríguez, *Relevancia normativa en la justificación de las decisiones judiciales*, 98–9; J.L. Rodríguez and G. Sucar, 'Las trampas de la derrotabilidad. Niveles de análisis de la indeterminación del derecho', 119–20; J.C. Bayón, 'Derrotabilidad, indeterminación del derecho y positivismo jurídico' in J.C. Bayón and J.L. Rodríguez, *Relevancia normativa en la justificación de las decisiones judiciales*, 182–5.

[3] R. Guastini, 'Defeasibility, Axiological Gaps, and Interpretation'.

on—would beg too many questions. Specifically, I am not going to assume that the norms I shall be discussing belong to a normative system, defined by membership criteria, and perhaps including priority rules among conflicting norms. This would to a large extent trivialize the issue and, by the way, when dealing with law or morality, this assumption is, to say the least, troublesome. There is, however, one tenet which is itself deeply controversial and is crucial to my whole enterprise. I shall assume that at least some normative issues (i.e. issues as to whether, for a given antecedent, a normative consequence follows) have a right solution—that there is a right answer to such issues in at least some cases. Nothing will be said to justify this assumption.

B. Specificationism: its promises . . .

The main phenomenon underlying talk of defeasible norms, the main ground for claiming that some norms are defeasible, is the possibility of conflicts between norms. According to an intuitionist model of the resolution of norm conflicts, what we have to do in such cases is balance the conflicting norms, in order to see which one, in the case at hand, prevails. (What this might mean is, of course, rather obscure.) Specificationism recommends a purportedly different course. When facing a conflict—or, when application of the relevant norm to a given case leads to a verdict we deem inappropriate, or unsatisfactory (this, of course, involves a value judgement)[4]—what we have to do is specify (that is, suitably restrict the domain of application of) at least one of the norms, or the relevant norm, so that, thanks to the inclusion of further conditions within its antecedent (thanks to '*glossing the determinables*')[5] the conflict—or the unsatisfactory verdict—eventually vanishes.[6]

This looks breathtakingly simple, and very promising. Obviously, a sensible specification will have to satisfy some constraints, excluding trivializing or irrelevant moves.[7] But, quite apart from the merits of these requirements, which I am not going to discuss here, the appeal of specificationism lies in what it promises. Specificationism claims to

[4] Whether or not these two hypotheses are in fact different I shall not consider here.

[5] Specification 'narrow[s] a norm by adding clauses spelling out where, when, why, how, by what means, to whom, or by whom the action is to be done or avoided'. See H.S. Richardson, 'Specifying, Balancing, and Interpreting Bioethical Principles', *Journal of Medicine and Philosophy*, 25 (2000), 289.

[6] H.S. Richardson, 'Specifying Norms as a Way to Resolve Concrete Ethical Problems', *Philosophy and Public Affairs*, 19 (1990); H.S. Richardson, *Practical Reasoning about Final Ends* (Cambridge: Cambridge University Press, 1994), ch. 4; H.S. Richardson, 'Specifying, Balancing, and Interpreting Bioethical Principles'; T.M. Scanlon, 'Intention and Permissibility', *Proceedings of the Aristotelian Society*, Suppl. vol. 74 (2000), but see also T.M. Scanlon, *What We Owe to Each Other* (Cambridge, Mass.: Harvard University Press, 1998), 197–202; and T.M. Scanlon, 'Adjusting Rights and Balancing Values', *Fordham Law Review*, 72 (2003), 1478, the latter on the related—but different—issue of 'adjusting', or specifying, institutionally defined rights. For a similar proposal, relying on Alchourrón's understanding of defeasible conditionals, see J.J. Moreso, 'Conflitti fra principi costituzionali', *Ragion Pratica*, 18 (2002); and, for a somewhat hedged version of the claim (which will not be discussed here), J.J. Moreso, 'Cristina Redondo sobre Razones y normas' *Discusiones* 4 (2005) in J.L. Rodríguez (ed.), 'Razones y normas', sect. 4. A parallel suggestion concerning conflicts between moral rights is found in R. Shafer-Landau, 'Specifying Absolute Rights', *Arizona Law Review*, 37 (1995).

[7] H.S. Richardson, 'Specifying Norms as a Way to Resolve Concrete Ethical Problems', 295–7; H.S. Richardson, *Practical Reasoning about Final Ends*, 72–4. These are mainly informal constraints. Formal constraints on specification should parallel the set of formal conditions which, according to Alchourrón, define a revision function for defeasible conditionals. See C.E. Alchourrón, 'Detachment and Defeasibility in Deontic Logic', *Studia Logica*, 57 (1996), 5–18; C.E. Alchourrón, 'On Law and Logic', in this volume, ch. 2; C.E. Alchourrón, 'Para una lógica de las razones prima facie', *Análisis filosófico*, 16 (1996), 113–125. For our purposes, it is not necessary to dwell on these matters here.

be a third way beyond balancing and subsumption, superseding both of them.[8] Thanks to specification, it is claimed, we can overcome conflict, and inappropriate verdicts, through a deeper, more adequate, understanding of the relevant norms.[9] This ensures real progress; we are not merely changing our minds as to what the norms are that we should follow, rather we are refining, qualifying them, on the basis of a deeper understanding of their content.[10] This, it is claimed, is something wholly different from an intuitive balancing of conflicting requirements—an altogether arbitrary procedure, leaving us stuck with abstract, unrefined requirements, liable to come into conflict over and over again. Specification is a reasoned way: the conditions we include within the antecedent of our norms, restricting their scope, count as reasons—reasons why a conflict should be resolved one way or the other, or a novel case distinguished from previous ones, and an inappropriate verdict avoided. Thus, specification brings about an enrichment of our normative outlook, leading us to more finely grained distinctions, refining our normative sensitivity, and our ability to discriminate in matters normative. In order to make specification possible, of course, we have to read our initial norms as defeasible (qualifications such as 'generally speaking', or 'for the most part', 'in most cases', etc. will have to be prefixed to our normative conditionals).[11] They will have to be understood as including *ceteris paribus* clauses: refinement and qualification deploy within the logical space left open by such clauses, sometimes defusing the link between antecedent and consequence. Specificationism is, thus, a species of defeasibilism.

We have to distinguish two main varieties of the specificationist strategy, one claiming that, through specification, we get to contributory norms (i.e. norms stating *pro tanto* reasons), the other claiming that specification leads us from contributory to 'all-things-considered' norms (or, to 'overall' reasons).[12] The former is most plausibly understood as the view that, thanks to specification we move from less specific to more specific contributory norms.

[8] H.S. Richardson, 'Specifying Norms as a Way to Resolve Concrete Ethical Problems', 279–80: 'a third, more effective alternative'; 'a third . . . operation'; 'a true third way, rather than just a mixture'; H.S. Richardson, 'Specifying, Balancing, and Interpreting Bioethical Principles', *passim*.

[9] The model of specification makes us understand 'how our ethical precepts, many of which are very general and abstract, can reach concrete cases without generating unacceptable conclusions'. H.S. Richardson, 'Specifying Norms as a Way to Resolve Concrete Ethical Problems', 284. I remind the reader (see section A) that my topic is not, specifically, *moral* norms. But there is nothing in the model of specification limiting its scope to the moral domain.

[10] Resort to specification 'helps ensure . . . that the reasonable motivation behind the initial, unqualified norm is still captured by what one ends up with'. Ibid., 284. Thus, the notion of specification 'can . . . explain how a moral theory can remain the subject of a more or less stable attachment despite the sort of revision that moral conflicts engender' (ibid.). Specification affords 'the progressive refinement of a theory that remains the same in essentials' (ibid., 298; see also H.S. Richardson, *Practical Reasoning about Final Ends*, 171); 'the model of specification learns from the conflicts it faces, exploiting their friction to push off toward a more concrete and definite understanding of the relevant norms' (H.S. Richardson, 'Specifying Norms as a Way to Resolve Concrete Ethical Problems', 308). Specification is 'a relation between two norms that allows one to trace the significant continuities in a path of revision' (H.S. Richardson, *Practical Reasoning about Final Ends*, 245).

[11] H.S. Richardson, 'Specifying Norms as a Way to Resolve Concrete Ethical Problems', 292–3; H.S. Richardson, *Practical Reasoning about Final Ends*, 70–2; H.S. Richardson, 'Specifying, Balancing, and Interpreting Bioethical Principles', 305, fn. 8.

[12] The former is Richardson's version (cf. H.S. Richardson, 'Specifying Norms as a Way to Resolve Concrete Ethical Problems', 286–7, 294), the latter Scanlon's (though neither of them uses these phrases). The terms 'contributory' and 'overall', as used here, are drawn from J. Dancy, *Ethics Without Principles* (Oxford: Clarendon Press, 2004).

The identity assumption is part and parcel of specificationism, in both versions. The leading thought in specificationism, it seems to me, is that, in order to solve the normative puzzle we are facing (be it a conflict, or an inappropriate verdict), we have to work out what the relevant norm, or norms, properly understood, involve. What is required, in order to overcome the difficulty, is a better, deeper—more adequate and, therefore, more articulated—understanding of the norm, or of the conflicting norms.[13] The point is sometimes made by distinguishing between a norm, or principle, and its different—more or less adequate, as the case may be—'formulations'.[14] This is emphatically not the distinction between sentences and their meanings (see section A). Rather, a 'formulation' of a norm, in the now relevant sense, is to be understood as a way, one of the many possible ways, of grasping its supposedly complex and finely grained content. One and the same norm may have different formulations: some of them will be mere shorthand, others will capture some of its details. In the process of specification, it is assumed, we do not, in fact, modify the norm, but, rather, its initial formulation; what we are engaged in is amending, refining, the norm's formulation (i.e. our inchoate, partial grasp of its content), in order to put it in line with the norm itself. (The norm is inadequately captured by its current formulation; the latter may be quite simple, the former, understood as its 'underlying idea', is a complex matter.)[15] It is,

[13] Scanlon explicitly emphasizes the identity assumption. T.M. Scanlon, 'Intention and Permissibility', *Proceedings of the Aristotelian Society*, Suppl. vol. 74 (2000), 308, fn. 8. While on Richardson's account, he claims, specification involves 'modifying' conflicting principles, according to his (Scanlon's) view, 'only one principle need be involved' and 'the process [of specification] is typically one of figuring out what the principle requires rather than modifying it'. This, I think, does not do justice to Richardson, in whose view, too, specification amounts to deepening our understanding of our initial norms (see the passages quoted above, nn. 9 and 10). True, Richardson explicitly defines specification as a relation between *two* norms. See, H.S. Richardson, 'Specifying Norms as a Way to Resolve Concrete Ethical Problems', 295; H.S. Richardson, *Practical Reasoning about Final Ends*, 72; H.S. Richardson, 'Specifying, Balancing, and Interpreting Bioethical Principles', 298. But, and this is all that is needed for my argument, it is essential to his enterprise that specification be understood as a way of bringing *the initial norm* to bear on the case at hand. The model of specification he claims (H.S. Richardson, 'Specifying Norms as a Way to Resolve Concrete Ethical Problems',290) 'starts from [the] recognition [that our norms are subject to revision], but establishes a kind of constancy or stability . . . This stability is essential to the claim that the initial norms are in some way *brought to bear* on concrete cases by means of more specific norms.' Ibid., 291: specification 'licenses us to call a modification . . . of an original norm still in some significant sense the *same* norm that we started out with'; thanks to specification it is 'not a self-contradiction to speak of modifying a consideration so that *it* applies'; the model of specification lays down 'conditions on the relation between the initial norm or norms and their modifications that explain how the original norms *are* being respected' (all emphases are in the original; cf. also ibid., 297, and parallel passages in H.S. Richardson, *Practical Reasoning about Final Ends*, 170–1). We need specification 'both to allow the development of a stable moral theory and to give us some assurance that the commitment that underlies the initial norm is being appropriately honored'; specification 'lets [the contours of this initial commitment] emerge on reflection' (H.S. Richardson, 'Specifying Norms as a Way to Resolve Concrete Ethical Problems', 292). Specification, in sum, ensures 'stability in the course of revision', ibid. Reading the relevant norms as defeasible ones—i.e. as qualified by clauses such as 'generally speaking', and the like—allows us to understand 'how a norm can be seen as "the same" before and after revision' (Ibid., 293). The identity assumption is rightly envisaged as crucial to specificationist views in R.M. Hare, 'Comments', in D. Seanor and N. Fotion (eds), *Hare and Critics* (Oxford: Clarendon Press, 1988), 263. Richardson overtly acknowledges the 'challenge' posed by Hare ('the challenge I am now addressing of how a norm can be seen as "the same" before and after revision'); his own version of specificationism is supposed to be able to cope with it. H.S. Richardson, 'Specifying Norms as a Way to Resolve Concrete Ethical Problems', 293, fn. 30; H.S. Richardson, *Practical Reasoning about Final Ends*, 71, fn. 2.

[14] See e.g. T.M. Scanlon, *What We Owe to Each Other*, 199: 'succinct verbal formulations' as 'mere labels for much more complex ideas'. A similar idea is mentioned in P. Väyrynen, 'Moral Generalism. Enjoy in Moderation', *Ethics*, 116 (2006), 725.

[15] The phrase is Scanlon's. See T.M. Scanlon, *What We Owe to Each Other*, 199.

thus, *the same norm* that we are supposed to be applying to the case, although, now, better understood, and in a suitably refined formulation.

The identity assumption is what underlies the idea, mentioned a few paragraphs above, that specificationism, in contrast to balancing, ensures (reasoned) progress. Specification does not amount, trivially, to changing our minds as to what the norms are that we should follow. Rather, it brings about a fuller articulation of our normative convictions.

C. ...and its shortcomings

There is, it seems to me, a simple argument that shows why this strategy fails. I shall call it 'the argument from negative conditions'. It runs as follows.

Our starting point is a norm claiming that when p is the case, C follows. Now, suppose that—as specificationism recommends us to do in the face of conflict—we grant, however, that when p and q is the case, C does not follow. (So that we now acknowledge, *ex post*, that our starting point had to be read as 'if p and not q, then C'.) And we further grant—always following specificationist advice—that, when p, q, and z are the case, C does indeed follow. (So that we now acknowledge, *ex post*, that our previously amended norm had to be read as 'if p and q and not z, C does not follow'). And so on, each time allowing, along specificationist lines, that the addition of a further conjunct may reverse the verdict, i.e. may now show that our previous verdict was right under a hitherto unstated negative condition. My question is, unless we may reasonably conceive of an exhaustive inventory, or list, of all possible conditions p, q, z, w etc. (or properties P, Q, Z, W etc.) which may obtain, and which may make a difference as to whether C follows, and of a full specification of the ways in which such conditions would make a difference (i.e. unless we deem conceivable, and at least in principle epistemically accessible, an exhaustive, ultimate list of all potentially normatively relevant properties or conditions, defining in advance the universe of all possible cases), what sense can be made of saying that, in proceeding as we do, and as specificationism recommends us (thus, in granting exceptions, exceptions to the exceptions, and so on), we are nonetheless going on in specifying—revising, refining, qualifying, amending—one and the same norm? No sense at all, it seems to me.

Further, is the notion of such an inventory and specification a coherent notion at all? No, I think.[16] For two reasons: (1) new, potentially normatively relevant properties may come into existence; and (2) the history of human interactions is intentionally laden.

The first of these two points is, I think, intuitive. (Consider the property of something being a microchip, a nation state, or an academic.) Besides, the meaning of old words may change, so that they come to designate new things. And further, if we allow thick concepts in the antecedent of norms (and I cannot see why we shouldn't, as a matter of grammar, I mean), the point becomes even more obvious.

As to the second point, by saying that the history of human interactions is intentionally laden, and that this makes the notion of a list such as the one envisaged above incoherent, I have in mind two considerations. (1) Practical problems typically arise in

[16] Cf. for a related point J. Dancy, *Moral Reasons* (Oxford: Blackwell, 1993), 67. In the terminology of Alchourrón and Bulygin, such a list may be termed an 'ultimate (normative) thesis of relevance (or hypothesis)'. C.E. Alchourrón and E. Bulygin, *Normative Systems* (Vienna: Springer, 1971), 103–6. (There is no difference, in the present context, between 'thesis of relevance' and 'relevance hypothesis'; ibid., 106. On this last point, see also M.C. Redondo, 'Reglas 'genuinas' y positivismo jurídico' in P. Comanducci and R. Guastini (eds), *Analisi e diritto 1998* (Turin: Giappichelli, 1999), 256.

the course of interactions between a plurality of agents; (2) an action may be described in a plurality of ways, and it is conceived—by the agent herself, or by others—under one or more such descriptions (not necessarily the same).[17] Both phenomena stem from the fact that the mind is endowed with intentionality. I shall consider them in turn.

(1) Practical problems (i.e. problems as to what is to be done in given circumstances) typically involve many individuals. Which properties are relevant in order to determine how one should act, may depend, and often does depend, on the past interaction of these individuals.[18] Often what the present case is, what its normatively relevant features are, does not boil down to what can be read off from its present configuration, but also depends on its history—i.e. on the path along which the relevant individuals have reached the present situation.[19] (Whether the fact that A promised B to phy is, now, a reason for phy-ing may depend on how, in the past, B behaved toward A—or maybe toward C, with whom A, though not knowing him personally, shares some important convictions, such as convictions as to whether, and when, promises should be kept.) The set of all possible combinations of potentially normatively relevant considerations, when account is taken of the path-dependence of cases, cannot be exhausted. This should not be understood in the sense that, due to its complexity, an extremely powerful computer would be needed in order to articulate all relevant possibilities. Rather, this set cannot in principle be exhausted. Human agents have intentionality; intentional states of increasing complexity (intentional states having as their objects other intentional states, and so on), including sets of (both synchronically and diachronically) interlocking intentional states of different levels (e.g. mutual beliefs or common knowledge) may be—and often are—relevant as to what the shape of interaction, and its history, is. (So, for example, A's beliefs about B's, or others', intentions or beliefs, or his hopes about others' future attitudes and behaviour, and so on, may be relevant.) Relations between such states may both be internal—i.e. relating to their content—or causal. (It may be relevant whether B expects A to keep his promise, and this may perhaps depend on what he expects A to expect from him, B, as a consequence of the way in which, as B—perhaps mistakenly—believes, A behaved in a past, similar case; but, maybe, A then behaved as he did in order to induce B to expect him, now, to expect a given behaviour from B, and so on.)

(2) To this it should be added that an action may be conceived, by the agent herself or by others, under an indefinite plurality of descriptions, often different for different individuals.[20] (The agent may not see that what he is doing falls under a certain description, something which may be apparent to some other individual.) There is, it seems to me, no reason to believe that the vocabulary of possible action-descriptions is

[17] G.E.M. Anscombe, *Intention* (Oxford: Blackwell, 1957, 1972), 11.

[18] Needless to say, the argument that follows concerns features which may turn out to be *normatively* relevant. (This is what the list which is at issue in the text is all about, of course.) This dispels a serious misunderstanding, grounding a purported objection against a previous version of the argument (M.C. Redondo, 'Razones jurídicas, Respuesta a Caracciolo, Celano y Moreso' in J.L. Rodríguez (ed.), 'Razones y normas', *Discusiones*, 4 (2005), 150ff.), replying to B. Celano, 'Podemos eligir entre particularismo y universalismo?', in J. Rodríguez (ed.), 'Razones y normas'.

[19] This is the main root for the distinctively particularist idea that there is a narrative dimension in practical rationality.

[20] This point was made, in this connection, by J.J. Moreso, 'Dos concepciones de la aplicación de las normas de derechos fundamentales' in J. Beregón, F. Laporta, J.R. de Páramo, and L. Prieto Sanchís (eds), *Constitución y derechos fundamentales* (Madrid: Centro de Estudios Políticos y Constitucionales, 2004), 491; B. Celano, 'Podemos eligir entre particularismo y universalismo?', 113.

finite. And, obviously, which [illegible] normatively relevant.

These two factors combine: the [illegible] interaction between a plurality of [illegible] tional states (of increasing [illegible] descriptions of the relevant [illegible] known (or, they are expected to be [illegible] known, and so on) to some of the [illegible] Under which description a given [illegible] it to fall according to B's lights, etc.) [illegible] what led to the present situation [illegible] descriptions, and of possible interactions [illegible] mines (just as it may sometimes make [illegible] shape of the case at hand is.

In sum, human interaction brings [illegible] normatively relevant properties, and [illegible] This multiplicity is non-computable [illegible] tively relevant properties, or kinds of [illegible] of all possible novel plots is.

Thus, the claim that, in following [illegible] the same norm proves hollow. It would [illegible] have to be *absent* for the norm to [illegible] Negative conditions cannot be exhaustive [illegible]

The point is not, it should be noted, [illegible] relevant negative conditions are [illegible] sensibly as the case may be, we [illegible]

[21] The argument from negative co[illegible] according to him, a sound) argum[illegible] R. Holton, 'Principles and Particula[illegible] (2002), 196–7. Holton notes that the [illegible] the norm under consideration, a 'tha[illegible] properties, apart from those already me[illegible] and that no other principle is, either).' [illegible] such a norm to a given case, a normati[illegible] This is, I think, correct. The trouble a[illegible] antecedent of the norm amounts to th[illegible] addition of 'that's it' as a premise in o[illegible] precisely, where the argument from n[illegible] noticed that Holton explicitly acknow[illegible] clauses in the antecedent of norms am[illegible] related argument J. Hage, 'Law and [illegible] Hints to the argument from negative c[illegible] *derecho: Deber jurídico y razones para* [illegible] 346–51 (relating to moral norms only) [illegible] in Legal Reasoning' in Z. Bankowski [illegible] (Dordrecht: Kluwer, 1995), 143; [illegible] B. Hooker and M. Little (eds), *M*[illegible] H. Prakken and G. Sartor, 'The Thi[illegible] (the 'qualification problem'); P. Väy[illegible] (2006), 725, 736. Its ancestors are in[illegible] A. Flew (ed.), *Logic and Language* (1[illegible] *Concept of Law* (Oxford, Clarendon P[illegible] in J.C. Bayón, 'Derrotabilidad, indete[illegible]

finite. And, obviously, which description is picked out by whom, may turn out to be normatively relevant.

These two factors combine. The history of the present case—the course of interaction between a plurality of individuals—is constituted, *inter alia*, by intentional states (of increasing complexity, etc.); their contents comprise different descriptions of the relevant actions, descriptions under which these actions are known (or, they are expected to be known, or are hoped to be expected to be known, and so on) to some of the agents involved, but perhaps not to others. Under which description a given action falls according to B's lights (or, A believes it to fall according to B's lights, etc.) may matter, as far as the narrative concerning what led to the present situation is concerned. Thus, this plurality of possible descriptions, and of possible intentional states having them as their content, determines (just as it may sometimes make indeterminate) what the normatively relevant shape of the case at hand is.

In sum, human interaction brings about an indefinite plurality of potentially normatively relevant properties, and of relevant histories in which they are instantiated. This multiplicity is non-computable. The notion of the set of all potentially normatively relevant properties, or kinds of cases, is misconceived, just as the notion of the set of all possible novel plots is.

Thus, the claim that, in following specificationist advice, we are specifying one and the same norm proves hollow. It would have some bite, if we could tell what properties have to be *absent* for the norm consequence to follow. This is, however, impossible. Negative conditions cannot be exhaustively enumerated.[21]

The point is not, it should be noticed, that we can never legitimately assume that all relevant negative conditions are satisfied. On the contrary, this is what, more or less sensibly as the case may be, we do all the time. What I am claiming is, rather, that

[21] The argument from negative conditions is close to what Holton takes to be the main (and, according to him, a sound) argument supporting particularism, the 'supersession argument'. R. Holton, 'Principles and Particularism', *Proceedings of the Aristotelian Society*, Suppl. vol. 76 (2002), 196–7. Holton notes that the argument may be blocked by including, in the antecedent of the norm under consideration, a 'that's it' condition (stating that in the case at hand no further properties, apart from those already mentioned in the antecedent of the norm, are normatively relevant, and that no other principle is, either). Thus, by adding 'that's it' as a premise to an inference applying such a norm to a given case, a normative solution for the case will deductively follow (ibid., 199–200). This is, I think, correct. The trouble arises when we notice that inclusion of the 'that's it' clause in the antecedent of the norm amounts to the requirement that all negative conditions be satisfied; and that addition of 'that's it' as a premise in our inference amounts to asserting that they are satisfied. This is, precisely, where the argument from negative conditions shows the main difficulty to be. (It should be noticed that Holton explicitly acknowledges, in the *Abstract* of his paper, that inclusion of 'that's it' clauses in the antecedent of norms amounts to reading them as holding only *ceteris paribus*.) Cf. for a related argument J. Hage, 'Law and Defeasibility', *Artificial Intelligence and Law*, 11 (2003), 237. Hints to the argument from negative conditions may be found in, e.g. J.C. Bayón, *La normatividad del derecho: Deber jurídico y razones para la acción* (Madrid: Centro de estudios constitucionales, 1991) 346–51 (relating to moral norms only); J. Dancy, *Moral Reasons*, 57, 77, 80–1; G. Sartor, 'Defeasibility in Legal Reasoning' in Z. Bankowski et al. (eds), *Informatics and the Foundations of Legal Reasoning* (Dordrecht: Kluwer, 1995), 143; D. McNaughton and P. Rawling, 'Unprincipled Ethics' in B. Hooker and M. Little (eds), *Moral Particularism* (Oxford: Clarendon Press, 2000), 262–3; H. Prakken and G. Sartor, 'The Three Faces of Defeasibility in Law', *Ratio Juris*, 17 (2004), 120 (the 'qualification problem'); P. Väyrynen, 'Moral Generalism. Enjoy in Moderation', *Ethics*, 116 (2006), 725, 736. Its ancestors are in H.L.A. Hart, 'The Ascription of Responsibility and Rights' in A. Flew (ed.), *Logic and Language* (1st series, Oxford: Blackwell, 1951), 147–8; H.L.A. Hart, *The Concept of Law* (Oxford, Clarendon Press, 1961), 125–6, 131,as these passages are illuminatingly read in J.C. Bayón, 'Derrotabilidad, indeterminación del derecho y positivismo jurídico'.

specificationism requires, in principle, listing all negative conditions. And this, precisely, is what cannot be done.

In other words, specificationism, as a strategy for dealing with norm conflicts or inappropriate normative verdicts, is engaged in what M.N. Lance and M.O. Little aptly call the 'usual quest' of theory, 'which is to spend all our time filling in the holes in our generalizations'.[22] The quest proves, as they claim, 'deeply misguided',[23] not because the process never ends (mine, here, is not a claim of the 'we cannot go on ad infinitum' kind), but because it does not even get started in the first place. Negations cannot be enumerated. Thus, all (justified) verdicts turn out to be particular.

The argument from negative conditions leads to the rejection of both versions of specificationism, the contributory and the 'all things considered' ones. As regards the latter, the point is obvious. Achieving a fully specified, 'all things considered' norm, thereby ruling out the possibility of further, unspecified exceptions (apart from those already built into the norm itself) would require us to be in a position to draw a list of all potentially relevant properties of the kind mentioned. And this, we have seen, is misconceived. The former version, too, proves untenable although for a different reason. If what we achieve through specification is still a *pro tanto* reason, it cannot be excluded that the amended norm will itself come into conflict with another norm, thereby showing itself in need of further specification. But if further specification may lead to an inversion, in the present case, of the previously issued verdict (thereby showing, *ex post*, that it was justified only under an hitherto unspecified negative condition), and if—as this version of the strategy maintains—this possibility always remains open (so that we are forced to conclude that our verdicts are justified, whenever they are, under an indefinite and unspecifiable set of negative conditions), then no progress at all is made, it seems, through specification. We were, and still remain, at sea.

I am emphatically not *assuming*, as a (dubiously) self-evident postulate of a sort, that in order to make any progress, we have to be able to get from the contributory to the 'all things considered'. Rather, this is what the argument from negative conditions purports to *show*, that, unless this proves to be a viable move, justified verdicts always hold only under an unspecifiable set of negative conditions. And the move, as we have seen, is not available. To repeat, the worry is not that we cannot go on ad infinitum. This is what we actually do. The problem is, rather, that the supposed progress assured by specification doesn't even get started. Specification does not provide—nor does it purport to provide—*ex ante* guidance as to how our starting points have to be refined (this is a substantive matter).[24] But neither does it provide—nor can it purport to provide—guidance *ex post*, as to how a certain kind of case has from now on to be decided (further specification may, *ex hypothesi*, be called for, thereby reversing our previous verdict, i.e. showing it as holding only under an hitherto unspecified negative condition). So what kind of guidance does the norm (be it in its still-to-be-amended, or its already-amended guise) provide? Once again, all justified verdicts turn out to be particular. Specification only engenders an appearance of progress, where none is made.[25]

[22] M.N. Lance and M.O. Little, 'Defeasibility and the Normative Grasp of Context', *Erkenntnis*, 61 (2004), 453. Lance and Little are addressing, here, 'epistemic theory'.

[23] Ibid.

[24] H.S. Richardson, 'Specifying, Balancing, and Interpreting Bioethical Principles', 302, 305, fn. 4.

[25] I should stress that my argument says nothing about the merits of specificationism (of the contributory variety) as a decision-making procedure in institutional settings, e.g. public policy choices by administrative agencies, or in courts. Specification is advocated as a mode of public deliberation by administrative bureaucracies in H.S. Richardson, *Democratic Autonomy* (Oxford: Oxford University Press, 2002), 104, 227–8, 235, 237–9. Just as, even granted that practical reasoning is basically

All this mandates, I submit, our rejection of the identity assumption. Or, rather, it shows that the first way of construing it is a blind alley.

In spelling out the identity assumption we may wish to say that recalcitrant cases may be accommodated either because they are already provided for, albeit implicitly, by the norm, or because the norm may be refined so as to allow for the recalcitrant case as an exception ('if p then q, except when p and z'). So understood, recalcitrant cases qualify only as prima facie exceptions. On close inspection, or after suitable revision, they turn out to be provided for in the antecedent of the norm. Such a way of dealing with recalcitrant cases does not, I will say, treat them as true (as opposed to merely prima facie) exceptions. It is this way of construing the identity assumption that the argument from negative conditions shows to be hollow. Why not simply say that, in dealing with our normative issue (be it a norm conflict, or a verdict which we deem inappropriate), we are now simply changing our minds, i.e. we are now realizing that the norm we had endeavoured to apply to the case at hand was wrong (or that it was no norm at all), and substituting it with a different one? In fact, when we incorporate, along these lines, an exception in the antecedent of a norm, 'there is no difference between adding an exception to a rule and simply changing it'.[26] Moreover, where does the supposedly deep, qualitative, difference from intuitive balancing of conflicting considerations lie? When, facing a norm conflict, we have to decide which of the two norms is to be modified (in fact, discarded, and substituted with a different, more specific, one), and how or when, facing an inappropriate verdict, we have to decide how the norm's antecedent is to be modified, so that the case at hand does not fall under it any more, what we have to do is, it seems, just balance conflicting normative considerations.[27]

particularist, there may be good reasons for imposing a rule-based decision-making procedure (decision-making by entrenched prescriptive generalizations) on certain classes of choosers, or in certain decisional environments (F. Schauer, *Playing by the Rules. A Philosophical Examination of Rule-Based Decision-Making in Law and in Life* (Oxford: Clarendon Press, 1991), ch. 7; see also F. Schauer, 'On the Supposed Defeasibility of Legal Rules' in M.D.A. Freeman (ed.), *Current Legal Problems*, 51 (1998), 233, 237), so there may be good reasons for proceeding by specification in some institutional contexts.

[26] F. Schauer, 'Exceptions', *The University of Chicago Law Review*, 58 (1991), 893: 'the logical emptiness of the idea of an exception as an analytically distinct concept'; cf. also ibid., 873, and R.H. S. Tur, 'Defeasibilism', OJLS, 21 (2001), 359. This is the root of Alchourrón's dissatisfaction with justification on the basis of defeasible conditionals: so understood, resort to defeasible conditionals is, in fact, belief revision in disguise (on Alchourrón's views about this issue see J.L. Rodríguez and G. Sucar, 'Las trampas de la derrotabilidad. Niveles de análisis de la indeterminación del derecho', 105); J.C. Bayón, 'Por qué es derrotable el razonamiento jurídico?', 265.

[27] This, of course, also depends on how 'balancing' itself is understood. If by 'balancing conflicting considerations' what is meant is staring at the case at hand, and solemnly declaring 'here, N1 overrides N2', period, then of course this would not be a sensible way of proceeding. But this is not a fair picture of what balancing may be taken to be, either. Balancing, as such, is compatible with—and sensible balancing certainly involves—explaining why, as one sees things, in such and such a case N1 should be taken to prevail over N2, i.e. giving reasons as to why, in this kind of case, the former overrides the latter. This, to repeat, is what we do all the time and, as a result, what we get are more specific norms. What is at issue in the text is specificationism's claim to be providing something qualitatively different from this, and to be assuring a kind of progress, of enrichment and rational articulation in our normative outlook, which balancing would be incapable of bringing about. Intuitive balancing, Richardson argues (H.S. Richardson, 'Specifying Norms as a Way to Resolve Concrete Ethical Problems', 283; cf. also ibid., 287–8), lacks 'discursive rationality'. 'The question for balancers' is 'how their weightings have to be explained or justified' (ibid., 282). And, Richardson claims (ibid., 282–3), 'to the extent that the balancing is genuinely distinct from application [i.e. subsumption; ibid., 281] it affords no claim to rationality, for to that extent its weightings are purely intuitive, and therefore lack discursively expressible justification'. The same, however, may be said about the choice, in the face of a conflict, as to which norm to specify, and how. In deciding whether to specify a norm, or which one to specify, and how, what we are doing is, trivially, balancing conflicting considerations (and this,

Or at any rate specificationism tells us nothing different, or more, than this.[28]

Thus, on this reading, the identity assumption proves hollow. Talk of defeasible norms is mere rhetoric. Exceptions are treated as only prima facie ones: they are, in fact, incorporated in the antecedent of different norms. But, does the notion of a defeasible norm allowing for true exceptions make sense? I shall discuss this issue below.

D. From defeasibility to particularism

The argument in the previous section concerned specificationism. Specificationism claims to provide a way of dealing with norm conflicts which is consistent with the identity assumption. This is where its main attraction lies. It has been shown, however, that, precisely on this account, specificationism fails. On specificationist premises, the identity assumption withers away.

This, by itself, does not prove that defeasibility claims in general are incapable of living up to what they promise: that they necessarily fall short of warranting, as they by stipulation (see section A) imply, the identity assumption. Arguments analogous to the one deployed in the previous section can, however, be raised against other versions of defeasibilism. The way specificationism falls short of its promise is paradigmatic, I suggest, of the failures of defeasibilism, in many of its forms.[29]

To illustrate this point let us briefly consider a second, similar way of substantiating the identity assumption, already hinted at in the previous section. This is the claim that *implicitly* exceptions are already provided for by the norm. (This, it seems to me, is what is implied when it is said, as it is often the case, that defeasible norms are norms having implicit exceptions.) The norm's apparent formulation has to be understood, on this reading, as shorthand for a suitably complex counterpart. Its antecedent already contains, albeit in implicit form, the required exceptions. (How so? Maybe because it was the lawgiver's real or counterfactual intention that the norm should not hold for

for all we know, on specificationism's own lights). It is only when balancing is understood according to the (unfair) picture indicated above that the impression arises that specification ensures us something more than that.

[28] According to Richardson, one of the main features of the model of specification, distinguishing it from balancing, is that it goes hand in hand with justification, understood in coherentist terms. Specification can—and should—be practised as a way of improving the coherence and mutual support of our normative convictions (ibid., 300–2, 395, 397; H.S. Richardson, *Practical Reasoning about Final Ends*, 174ff., 185; H.S. Richardson, 'Specifying, Balancing, and Interpreting Bioethical Principles', 302; H.S. Richardson, *Democratic Autonomy*, 110). I have no quarrel with this, provided it is recognized that (1) if we allow (as this version of specificationism does indeed allow, and as the argument from negative conditions demands) for the standing possibility of exceptions to exceptions—or, specifications of specifications, revisions of revisions—coherence and mutual support turn out to be very weak constraints indeed; (2) this feature may be taken to differentiate specification from balancing only to the extent that the latter is given the (unfair) representation mentioned in the previous footnote.

[29] Of course, I have in mind here defeasibility of norms, as described in section A. Different notions of defeasibility, rooted in different concerns such as, e.g. the necessity of reasoning with incomplete information, or the proper ways of allocating the burden of proof, or the representation of knowledge in AI and law research; cf. G. Sartor, 'Defeasibility in Legal Reasoning', H. Prakken, *Logical Tools for Modelling Legal Argument. A Study of Defeasible Reasoning in Law* (Dordrecht: Kluwer, 1997), chs 1–5; J.C. Bayón, 'Por qué es derrotable el razonamiento jurídico?'; J. Hage, 'Law and Defeasibility'. H. Prakken and G. Sartor, 'The Three Faces of Defeasibility in Law' remain unaffected by this argument.

those cases, or for some other reason. Different explanations may be contrived.)[30] This is another main route to identity—thus, to defeasibility—claims.[31] And it, too, is shown to be a blind alley by the argument from negative conditions. Both strategies exhibit the same flaw: they engage us in the inane attempt at 'filling in the holes in our generalizations'. (Exceptions are not *real* holes, they in fact claim, either because they can be presently filled in or because they turn out to be already, albeit implicitly, filled in.)

Where does the failure of specificationism—or, if the suggestion at the beginning of this section is sound, of defeasibilism in many of its versions—leave us? Abandoning the identity assumption, while still allowing for the possibility of norm conflicts or inappropriate verdicts, amounts, I think, to endorsing a kind of particularism.

'Particularism' is an equivocal term, and many different positions may be labelled 'particularist'. This is not the place for sorting out varieties of particularism, their interrelationships, and so on. For our purposes, 'particularism' may be taken to mean what is entailed by the conclusion we reached, through the argument from negative conditions, in the previous section, namely that, when a specificationist strategy is followed, all justified normative verdicts turn out to be particular, i.e. they hold only under an indefinite set of negative conditions. Let us see what this conclusion amounts to.

What may be meant, here, by saying that justified verdicts remain particular is, in fact, two different things, leading to two different notions of particularism. I will list them both, since I believe there is no need, in the present context, to choose one of them and discard the other. The two claims are the following: (1) 'normative verdict particularism', namely, the claim that there are no absolute (i.e. 'all things considered', or overall) norms determining the deontic status of an act;[32] (2) the radical particularist claim ('normative valence particularism') that a feature that is a reason favouring an action in one case may be no reason at all, or even a reason against, in another case (i.e. features may shift their normative valence).[33] The two claims are independent of one another.[34] The two versions of particularism stemming from them are, first, normative

[30] None of them is endorsed here as a viable strategy. Whether such a claim can be plausibly made is a question I simply leave aside. (On what 'implicit' exceptions may be taken to be see J.L. Rodríguez, 'La derrotabilidad de las normas jurídicas', 94–101.)

[31] It was, I think, Alchourrón's way. Alchourrón's is a neat example of what Lance and Little call an 'enthymematic' strategy in dealing with recalcitrant cases (M.N. Lance and M.O. Little, 'Defeasibility and the Normative Grasp of Context', 438; see also M.O. Little, 'On Knowing the "Why". Particularism and Moral Theory', *The Hastings Center Reports*, 31 (2004), 37). It should be emphasized, however, that Alchourrón himself was well aware that this was an easy, but Pyrrhic victory over recalcitrant cases (see above, n. 26).

[32] Normative verdict particularism is the generalized form of 'moral verdict particularism', as defined in D. McNaughton and P. Rawling, 'Unprincipled Ethics', 258: 'the claim that there are no absolute principles for determining the overall deontic status of an act'. By 'particularism' is usually meant, in contemporary debates, moral particularism (that is, a doctrine about the behaviour of moral reasons). Since what we are concerned with here are normative structures in general (see section A), particularism will be recast, accordingly, as a doctrine about norms generally.

[33] 'Moral valence particularism': 'there are no properties, apart from the thin moral properties right, wrong, etc., that have universally and counterfactually invariant valence' (ibid., 268). This is, roughly, Dancy's position in *Moral Reasons* and *Ethics Without Principles*. I say 'roughly' because Dancy somehow allows for the possibility of invariant moral reasons, and does not straightforwardly deny that there can be true moral principles. But these complexities need not detain us here.

[34] D. McNaughton and P. Rawling, 'Unprincipled Ethics', 258.

verdict particularism as such; and, second, radical particularism: the conjunction of normative verdict particularism and of the radical particularist claim.[35]

Verdict particularism is, as it were, minimal particularism. It is compatible with there being a multiplicity of *pro tanto* reasons which may, and usually do, conflict with varying strength or weight (and no previously established priority rules, of course), and which on each case have to be balanced against one another (this is, roughly, Rossian intuitionism).

Understood as either claim (1), or as the conjunction of claims (1) and (2), particularism is a position allowing for the following claims: norms can and do in fact come into conflict; when a conflict arises, we cannot but strike a balance, declaring that, in such and such a case, norm 1 weighs more than norm 2—sometimes perhaps substituting N2 with another norm, N2*, more restricted in scope (specification)—the results of such balancing (and substitutions) are, however, open to revision when, as it is allowed, new conflicts arise; no revised norm may be held to be immune to further revision, thus we allow for the possibility that when properties P and Q are instantiated, N1 prevails over N2 (the latter perhaps being substituted by N2*), and that the possibility always remains open that there is a property Z such that, when P, Q, and Z obtain, N2 (or N2*) prevails over N1 (the latter being perhaps substituted by N1*) and so on. The argument from negative conditions entails particularism, in either of these two forms.[36]

Thus, the rejection (via the argument from negative conditions) of specificationism, and of relevantly similar versions of defeasibility claims, leads to particularism. The question is, does particularism (as defined above) amount to avowing that, in fact, no norm is being applied throughout? That, namely, we just see, case by case, what the right answer is? Is there any room left for reliable normative generalizations?

I shall take up this issue in section F. Before attempting an answer to this question, however, we need to pin down a point that emerged in the previous section.

E. Prima facie vs true exceptions

The argument so far has led us to a distinction between two different ways of talking about 'exceptions', in fact, two concepts of an exception.[37]

Exceptions are only prima facie if they may plausibly be understood as incorporated, maybe only implicitly (see section D) in the norms they are exceptions to, i.e. when we may plausibly claim that a proper understanding, or formulation, of the norm would envisage these kinds of cases: properly understood, or properly reformulated, the norm

[35] See ibid., 258–9.

[36] It should be noticed that particularism, as here understood, is not the claim that *individual* cases always escape, because of their inherent complexity or richness, the grip of conditional norms providing normative solutions for generic cases. (The notions of individual and generic case are defined in C. E. Alchourrón and E. Bulygin, *Normative Systems*, 27–30.) This claim is scarcely intelligible, and surely not a consequence of the argument from negative conditions. Rather, it is by virtue of *properties*, or, more generally speaking, features, serving as reasons, that new cases may be distinguished from previous ones—i.e. that a justified verdict may in the present case (a case sharing with the previous one the features that were sufficient reasons for the verdict there) be reversed (so that the previous verdict is now shown to hold under hitherto unspecified negative conditions).

[37] The distinction between these two concepts is somehow adumbrated in R. Dworkin, *Taking Rights Seriously* (London: Duckworth, 1977), 25. It is there entangled, however, in Dworkin's distinction between rules and principles.

provides for them so that, in resolving norm conflicts what we have to do is to ask ourselves 'how the relevant principles . . . are best understood'.[38]

When only prima facie, then, exceptions are nothing but further elements of the antecedent of a norm. That they are taken to be somehow implicit, or that they are formulated as separate items, or as a result of specification,[39] is only fortuitous—in principle, a proper understanding or formulation of the protasis of the norm could, and indeed would, lay them out on the same footing as other conditions.[40]

So, when is an exception a true exception? For a case to qualify as a true exception it must not be already provided for in a reasonably detailed and precise 'unless . . . ' clause attached to the norm. Two points need to be made here.

First, whether such a clause is conceived of as explicit or as merely implicit is immaterial here. When, and to the extent, it can truly and justifiably be claimed that an 'unless . . . ' clause of the required sort is implicitly attached to a given norm, the exception is only a prima facie exception, not a true one. (I am emphatically *not* assuming that such a claim can ever be truly and justifiably made. Maybe not. On this issue, I simply express no opinion.)

Second, the exception is not already provided for in *a reasonably detailed and precise* 'unless . . . ' clause. Obviously, open-ended or abstract clauses such as 'unless phy-ing would be unreasonable in the circumstances', or 'unless the circumstances demand otherwise', 'unless there are very good reasons for doing otherwise', etc. (e.g. promises should be kept 'at least in the absence of special justification'),[41] do not qualify as clauses of the required sort.[42] They can easily be given a particularist interpretation, as allowing for the possibility that some unspecified property, or set of properties, will

[38] T.M. Scanlon, 'Intention and Permissibility', *Proceedings of the Aristotelian Society*, Suppl. vol. 74 (2000), 309. See also ibid., 310: 'plausible moral principles do not merely state general requirements but also incorporate exceptions to these requirements'.

[39] Or even that they are characterized as *non-refutanda* (vs. *probanda*). G. Sartor, 'Defeasibility in Legal Reasoning', 1995), 131.

[40] F. Schauer, 'Exceptions', 872–3 and *passim*; see also F. Schauer, 'On the Supposed Defeasibility of Legal Rules', 227, 231; R. Dworkin, *Taking Rights Seriously*, 25; R.H.S. Tur, 'Defeasibilism', 359–60. In such cases, 'little more than deception is served by employing the language of exceptions'. F. Schauer, 'Exceptions', 895; see also ibid., 898–9. Prima facie exceptions may also be understood as the result of conflicts between non-defeasible norms, when a preference (or a hierarchy) between the conflicting norms is assumed (e.g. because a priority rule, such as *lex specialis*, is presupposed). See J.L. Rodríguez, 'La derrotabilidad de las normas jurídicas', 86; J.C. Bayón, 'Por qué es derrotable el razonamiento jurídico?', 287.

[41] Cf. T.M. Scanlon, *What We Owe to Each Other*, 200.

[42] Or, in other words, what would be required is a 'fairly definite and informative general account . . . of what count as potentially' relevant conditions. P. Väyrynen, 'Moral Generalism. Enjoy in Moderation', *Ethics*, 116 (2006), 737. Cf. G. Sartor, 'Defeasibility in Legal Reasoning', 143; R.H.S. Tur, 'Defeasibilism', 361–2: 'If A is, then B ought to be, unless there is an overriding reason to the contrary': 'one moves from a list of specific reasons, to a general catch-all residual category which unlike a list is not closed and cannot be closed'. A clarification is required here. There is, it seems to me (cf. J.C. Bayón, 'Por qué es derrotable el razonamiento jurídico?', 294), a difference between vagueness, as a feature of general terms, and the open-endedness of clauses such as those mentioned in the text. Contrary to Schauer (F. Schauer, 'On the Supposed Defeasibility of Legal Rules', 231), that exceptions covered by such clauses are 'not specified in advance' is not to be equated with the fact that—i.e. understood as meaning nothing but that—the extension of a vague term or phrase cannot be specified in advance. (This amounts to claiming that the distinction between the two kinds of 'unless . . . ' clauses in the text is not equivalent to, but rather cuts across, the distinction traced by Schauer (ibid., 231), between 'weak' and 'strong non-specifiability', according to whether exceptions can or cannot be specified in advance at least 'by type'. Here, 'type' blurs the important distinction between, on the one hand, the indeterminacy of the extension of ordinary general terms—i.e. vagueness—and, on the other hand, the indeterminacy of open-ended, abstract clauses.)

unexpectedly prove relevant in the circumstances. In order to work out whether such a clause applies to a given case, one has to work out whether applying the consequence would be justified with regard to a wide range of normative considerations. And this sort of judgement, particularists will claim, may, and should, be accounted for along particularist lines.[43]

When is an 'unless . . . ' clause *sufficiently* detailed and precise so as to warrant us in labelling the exception a merely prima facie, not a true one? This, itself, is a normative question, to be answered by working out what 'reasonably' (a *reasonably* detailed and precise 'unless . . . ' clause) amounts to in a given context. We may easily, as usual, give examples falling at the extremes. Often, however, there will be cases in which determining whether an exception is a true or a merely prima facie one will require us to make up our minds about substantive issues. This is the province of *determinatio* of open-ended or abstract clauses, about which I have nothing of interest to say. (I remind the reader that what my argument is about are norms, not norm formulations. Thus, the issues to be dealt with here are not issues of interpretation—remember the caveat spelt out in section A).

True exceptions, then, are not (not even in principle) specifiable, and enumerable, in advance. There are two ways of reading this: (1) E is a true exception if and only if it cannot be provided for in advance; (2) the list of possible exceptions cannot be exhausted. The relevant understanding is the latter. What now turns out to be an unexpected, unprovided-for exception may, from now on, become a settled one; what is left open is the possibility that an exception to this exception will present itself, or, generally, that further exceptions will have to be acknowledged.[44] (Does the former reading even make sense? The only hypothesis I can think of is that of not-yet-existing properties.)

To sum up then, true exceptions occupy a middle ground between two extremes. On the one hand are exceptions already provided for, be it explicitly or implicitly, in the norm through reasonably detailed and precise 'unless . . . ' clauses (i.e. prima facie exceptions). On the other hand is the bare fact of the non-application of a norm. We

[43] Scanlon's suggestion (*What We Owe to Each Other*, 200–1), that we may get to the relevant principle, piercing through its ordinary 'succinct verbal formulation', by considering the 'point' of (that is, by 'understanding the rationale' for) the general requirement the latter expresses (thereby determining whether the present, 'new and difficult', case should count as an exception to the requirement) does not, by itself, weigh in favour of generalism. Particularists will give a particularist account of precisely this way of proceeding. A sort of middle ground is occupied, here, by the notion of principles which relate non-moral features to moral ones (or, conditional norms linking factual antecedents to normative solutions), but do so only by virtue of evaluative or normative riders (in such cases, 'the list of conditions is not open-ended, and it is knowable in advance', but 'the conditions cannot be spelled out in purely non-moral terms'; D. McNaughton and P. Rawling, 'Unprincipled Ethics', 268–72. It all depends on what the relevant evaluative or normative riders are.

[44] It might be argued that true exceptions are what Hart was getting at (H.L.A. Hart, 'The Ascription of Responsibility and Rights', 147–50, and especially in *The Concept of Law*, 135–6: rules have exceptions 'incapable of exhaustive statement'), as these passages are read, and freed from some confusions, in F. Schauer, 'Exceptions', 896–7: 'the way in which rules can be overridden in particularly exigent circumstances and still be rules, even if it is impossible to predict or to specify in advance what those exigent circumstances will be'; and in J.C. Bayón, 'Derrotabilidad, indeterminación del derecho y positivismo jurídico', 163–4, 176–81. See also F. Schauer, 'On the Supposed Defeasibility of Legal Rules', 225: Hart 'made clear that the claim of defeasibility was more than the mere claim that rules can be defeated upon the occurrence of *specified* defeating conditions. Rather, to Hart it was the very *unspecifiability* of the defeating conditions, "the use of the word 'etcetera'," that explained the operation of legal rules.' What I call 'true' exceptions are also the main focus of R.H. S. Tur, 'Defeasibilism', 301ff.; R. Holton, 'Principles and Particularism', *Proceedings of the Aristotelian Society*, Suppl. vol. 76 (2002), see esp. 207, 209, '*exceptio probat regulam in casibus non exceptis*'.

should clearly distinguish, of course, two different issues.[45] The first is how a norm regulates a given case; the second is given that the case at hand falls under the antecedent of the norm, whether the norm is in fact applied or not by a given subject (e.g. a judge) on a given occasion, for whatever reason (because, for instance, the outcome looks unfair or unjust to him, or because he was bribed).[46] The bare fact that a norm is not applied by a given subject on a given occasion is, of course, 'no reason to regard it as defeasible':[47] the case at issue does not, for this reason alone, count as an exception (neither as a prima facie nor as a true one). The notion of a true exception is the notion of an unprovided-for case, C, falling under the antecedent of the norm, in which, nevertheless, the norm consequence does not follow because the norm holds in normal circumstances, or *ceteris paribus*, and the circumstances constituting C are not normal: other things are not equal. It is here, as we shall now see in more detail, that there is room left, in the particularist picture, for reliable, though defeasible, generalizations, and that a second way of understanding the identity assumption gains some plausibility.[48]

F. Default reasons, normal contexts (or, From particularism to defeasibility)

Particularism—especially, radical particularism (as defined in section D)—is a much debated position. One of the objections most often raised against radical particularism is that it does not account for, or even contradicts, an apparently non-dispensable intuition, or platitude, governing moral thought and judgement, namely, that certain features of situations (e.g. that a certain action would cause needless pain to an unwilling innocent) seem to be reasons, as it were, in themselves or in their own

[45] E. Bulygin, 'Review of Jaap Hage's Law and Defeasibility', 247.

[46] 'A judge may correctly identify the applicable legal rule and then decide not to apply it'. Ibid., 247.

[47] Ibid., 248.

[48] One related issue is whether *legal* norms, specifically, are defeasible. Bayón is of course right in pointing out that, when judges decide not to apply applicable legal norms because the outcome is morally unsatisfactory (law and morality being conceived, here, as different, discrete, normative systems), this, by itself, does not show that the relevant norm is defeasible (J.C. Bayón, 'Por qué es derrotable el razonamiento jurídico?', 293–4; see also J.L. Rodríguez and G. Sucar, 'Las trampas de la derrotabilidad. Niveles de análisis de la indeterminación del derecho', 121–2, 144; F. Schauer, 'On the Supposed Defeasibility of Legal Rules', 230; J.L. Rodríguez, 'Derrotabilidad e indeterminación del derecho', 229). The issue is, however, which is the best way of rationally reconstructing the form of legal norms, or of some of them—as defeasible conditionals, or as indefeasible ones. (Why on earth should it be assumed that the burden of proof is on defeasibilists? That, namely, *unless* positive law—explicitly, or perhaps even implicitly?—provides to the contrary, legal norms should be held to be indefeasible conditionals? And, why not read legal norms excluding unspecified exceptions as defeasible? The burden of proof seems to me to be evenly distributed here.) True, the bare fact that judges sometimes decide that the consequence should not be applied in a case which, as it seems, does indeed fall under the norm's antecedent does not answer this question (Bulygin's point in the text). But, on the other hand, the fact—if it is a fact—that true exceptions, not specified in advance (but for open-ended, abstract, or generic clauses), have long been recognized, and are still recognized, as possible (that exceptional circumstances are often recognized as justifying exceptions) in legal culture (e.g. by judges and jurists) in many, or most, or perhaps all legal systems, and that open-ended, abstract, generic clauses of the relevant sort are pervasive in many, most, perhaps all legal systems, may plausibly be taken to justify a defeasibilist reconstruction of at least some legal norms (see R.H.S. Tur, 'Defeasibilism', 307–8; J. Hage, 'Law and Defeasibility', 232–3)—or at least as shifting the burden of proof to the other foot.

right; that is, they seem to be endowed with moral relevance—specifically, with a given (positive or negative) normative valence—by virtue of what they are, not to be made thus relevant by other features of the situations in which they are present.[49] While, on the other hand, other features (e.g. shoelace colour) seem to draw whatever moral relevance they may happen to have in a given situation from other features of the situation. Blurring this difference, it is claimed, amounts to 'flattening the moral landscape'.[50]

This is a serious charge, one particularism should prove able to cope with. Even the most radical particularist should, I think, acknowledge that some reasons look like genuine *pro tanto* reasons: that some considerations seem to have normative relevance in their own right, and that they standardly count in favour of, or against, actions (no special explanation is called for when they do). The device J. Dancy has put forward in order to cope with this phenomenon is the notion of a default reason: some considerations arrive already 'switched on' as reasons in favour of, or against, an action so that when they do so count 'there is nothing to explain'; 'it is only when things are not as they are "default-set to be" that we begin to ask questions'.[51]

Default reasons are an especially tricky issue for radical particularists.[52] What is relevant to our present purposes, however, is a (relatively) uncontentious point. Even if we grant that reasons behave as radical particularists claim they do, even if we grant that all properties may shift normative valence according to context, some room has to be left for the notion of what standardly, or normally happens (as far as which considerations are reasons for which actions is concerned):[53] generalizations about what is a reason for what must be allowed, provided they are read as 'written with the standard case in view',[54] i.e. provided they are understood as relating to what is normally the case—and, thus, as defeasible generalizations.

I have tried to show that some ways of moulding the concept of defeasibility of norms end up in hollow rhetoric; changes in our normative outlook in the face of recalcitrant cases (i.e. norm conflicts, inappropriate verdicts) are all that is at issue. This line of thought leads, I have also claimed, to particularism: rejection of a specificationist understanding of the identity assumption, coupled with an awareness of the phenomena prompting defeasibility claims (that is, recalcitrant cases of the kinds mentioned)

[49] Cf. J. Dancy, 'What is Particularism in Ethics?', ('Cos'é il particolarismo in etica?', Italian transl.), *Ragion Pratica*, 26 (2006), 113–132, on the '"feature-placing" aspect of moral deliberation and reasoning', meaning by this 'the way in which a case can be made for or against an action by starting off with certain features that seem to have a relevance in advance of any consideration of the context'.

[50] S. McKeever and M. Ridge, 'Turning on Default Reasons', *Journal of Moral Philosophy*, 4 (2007). For a similar complaint see D. McNaughton and P. Rawling, 'Unprincipled Ethics', 268, 273. M.N. Lance and M.O. Little, 'Particularism and Anti-Theory' in D. Copp (ed.), *The Oxford Handbook of Ethical Theory* (Oxford: Oxford University Press, 2005), 583.

[51] J. Dancy, *Moral Reasons*, 103, 230; J. Dancy, *Ethics Without Principles*, 112; J. Dancy, 'Defending the Right', *Journal of Moral Philosophy*, 4 (2007), 89, from where the quoted passages are drawn.

[52] See S. McKeever and M. Ridge, 'Turning on Default Reasons'; and Dancy's partial retraction and defence in J. Dancy, 'Defending the Right', *Journal of Moral Philosophy*, 4 (2007), 89–92.

[53] Cf. D. McNaughton and P. Rawling, 'Unprincipled Ethics', 268. Simple normative verdict particularism (as contrasted with normative valence particularism; see section D) may, of course, allow for normative *pro tanto* generalizations (this is the distinctive feature of Ross-style intuitionism). So, when simple normative verdict particularism is concerned, the worry whether particularism allows for reliable normative generalizations is in fact out of place. Both the argument from negative conditions and the distinction between prima facie and true exceptions, although compatible with simple normative verdict particularism, seem to weigh in favour of a somehow stronger form of particularism. And this is the line followed in the text.

[54] The phrase is from ibid., 269.

leads, via the argument from negative conditions, to particularism. But, as we have just seen, radical particularism itself leaves—or at any rate should somehow leave—the possibility open of reliable, though defeasible, generalizations about what is a reason for what;[55] generalizations stating what is normally the case, as far as what is a reason for what is concerned. It is through this notion, normalcy, that we may capture, and make some sense of, the grain of truth in the claim that norms are defeasible.[56]

Let us say, then, that norms state what are *normally* reasons for or against certain actions, or certain normative consequences. Norms are defeasible conditionals liable to true exceptions, i.e. conditionals such that the consequence follows, when the antecedent is satisfied, under normal circumstances only.[57] The crux of the matter is, of course, how is the qualification 'normally' ('in normal circumstances', etc.), to be understood? Here, it seems, normalcy includes, but does not boil down to, the notion of what happens, or holds, 'in most circumstances'.

M.N. Lance and M.O. Little's notion of a defeasible generalization, as one resting on a '(normative) grasp of privileged conditions', hits, I think, the nail on the head.[58] This would be a way of having 'robustly explanatory', illuminating generalizations with true exceptions—or, *mutatis mutandis*, defeasible conditional norms. The problem, however, is still there: how is the relevant notion of 'privileged conditions'—or, what amounts to the same, 'normalcy'—to be understood? Clearly, it cannot be by arbitrary *fiat* that some conditions count as normal ones but, as Lance and Little rightly emphasize, it cannot be a merely statistical matter either.[59]

Lance and Little hint at the possibility of having 'a grasp of the shape of "privileged conditions"',[60] thus developing 'a skill at understanding and recognizing what is

[55] This move from particularism to defeasibility is the theme of M.N. Lance and M.O. Little, 'From Particularism to Defeasibility in Ethics' in Lance et al. (eds), *Challenging Moral Particularism*, (New York: Routledge, 2008). See also M.N. Lance and M.O. Little, 'Particularism and Anti-Theory', 588–91. A hint may also be found in J.L. Rodríguez, 'Introducción' in J.L. Rodríguez (ed.), 'Razones y normas', 22. As Rodríguez rightly points out, particularists will claim that an Alchourrónian revision function for a defeasible conditional is indeterminate (ibid., 24, 25). The notion of a defeasible conditional whose revision function is indeterminate captures part of what talk of generalizations having true exceptions is about. (According to particularists, that is, norms are to be understood as defeasible conditionals liable to true exceptions.) The remaining part is the idea of normalcy, as discussed in the text.

[56] Schauer grants that it is compatible with decision-making by genuine rules that these stay open to reconsideration in 'especially troublesome circumstances'. F. Schauer, 'On the Supposed Defeasibility of Legal Rules', 234. See also ibid., 238–9; and F. Schauer, *Playing by the Rules*, 98, fn. 26 on the 'presumptive force' of rules ('very good reasons', 'an elevated standard of defeat'). This allows for the version of the identity assumption now under consideration (ibid.). See also, in the same vein, F. Schauer, 'Exceptions', 897, about the 'standard of exigency'.

[57] Cf. R.H.S. Tur, 'Defeasibilism', 359; R. Holton, 'Principles and Particularism', 207. That defeasibility of norms should be understood in terms of 'normalcy' (i.e. defeasible norms as norms such that the consequence follows, when the antecedent is satisfied, under normal circumstances only) is often acknowledged, often in passing only. See e.g. G. Sartor, 'Defeasibility in Legal Reasoning', 123; J.L. Rodríguez, *Lógica de los sistemas jurídicos*, 356; H. Prakken and G. Sartor, 'The Three Faces of Defeasibility in Law', 120.

[58] M.N. Lance and M.O. Little, 'Defeasibility and the Normative Grasp of Context'; 'Particularism and Anti-Theory', 589; 'Defending Moral Particularism'; 'From Particularism to Defeasibility in Ethics'. A similar proposal has been put forward by Väyrynen in P. Väyrynen, 'Moral Generalism. Enjoy in Moderation'. I shall not discuss affinities and differences between them (see ibid., 727, fn. 51).

[59] M.N. Lance and M.O. Little, ('Particularism and Anti-Theory', 588; 'Defeasibility and the Normative Grasp of Context', 438, 441, 444, 445) also purport to show that privileged conditions (understood in the relevant sense) may be quite rare. I find their arguments on this score unconvincing.

[60] M.N. Lance and M.O. Little, 'Defeasibility and the Normative Grasp of Context', 452. See also M.N. Lance and M.O. Little, 'Particularism and Anti-Theory', 591.

deviant and normal, what paradigmatic and emendational, what conceptually prior or central'.[61] Such a skill, and its related object (the 'shape' of normalcy) remain, however, rather obscure. As far as I can see, there are two main problems involved.

(1) *Background and shape.* Normalcy, in the relevant sense, is an irreducibly contextual notion, in two respects.
   (a) It cannot be exhaustively spelt out what, in a given case, normal circumstances amount to. (This is the idea underlying the argument from negative conditions, and the point of contrasting true with prima facie exceptions).
   (b) Normal conditions are the context within which norms apply, i.e. they are the background against which normative conditionals properly work. It is only against this background that consequences follow from their antecedents.

This has to be understood by keeping in mind the background–shape relationship in *Gestalt* psychology. In each context, trying to specify what conditions are normal is just like trying to see the background as coming to the fore, showing its shape (just like defining the shape of the background). This cannot be done.

(2) *'Is' and 'ought'.* Normalcy, in the relevant sense, is supposed to occupy an intermediate ground between facts and norms, or between 'is' and 'ought'. The idea of normal conditions is not the idea of a norm, or set of norms (or of a set of facts satisfying them); but nor is it reducible to the notion of a mere regularity, although it includes that. Normalcy is neither rule, nor regularity, but rather something intermediate between the two. And this is awkward. Consider the following two statements:
   Most cases are normal.
   Most of the time (or, usually), things follow their normal course.
   The relevant notion of normalcy only makes sense if these can be understood as meaningful, non-tautological true statements. And I confess that I find this puzzling. Is it (non-tautologically) true that most cases are normal, that usually things follow their normal course? If the argument so far is right, it should be possible to make some sense of statements such as these.

Most of the time, then, things follow their normal course. And it is against this background (i.e. in normal cases) that normative conditionals work—that is, that normative consequences follow from their antecedents—and that true exceptions become possible.

G. Conclusion

I do not wish to claim that, whenever we face a norm conflict, or an inappropriate normative verdict for a given case, what we should do is declare that circumstances are abnormal and that, alas, true exceptions are a standing possibility. That would be ludicrous. Nor is this entailed by the argument so far. Obviously, there is room for specification of our norms, refining the universe of cases which their antecedents incorporate, and for building prima facie exceptions in the antecedents of such revised norms. What I am claiming is that, when this is what we are engaged in doing, talk of

[61] M.N. Lance and M.O. Little, 'Defeasibility and the Normative Grasp of Context', 453.

defeasible norms amounts to mere rhetoric. The identity assumption has no bite, here. We are just revising our normative convictions—changing our minds, more or less sensibly or coherently, as a consequence of balancing conflicting normative considerations.

This, however, should not be confused with a different phenomenon, the emergence of true exceptions to genuinely defeasible norms. This possibility, I have tried to show, is conceptually open. And it is here that the identity assumption can be made sense of, although now in a contextualist framework: defeasible norms remain in place (though their consequences do not follow there) in abnormal circumstances—but which circumstances count as 'abnormal' is a contextual matter. (Defining what is 'normal' would be like seeing the shape of the background.)

The identity assumption, on this reading, rests on a necessarily implicit understanding of what counts as 'normal' conditions, i.e. conditions that usually fit our usual expectations. Most of the time, things follow their normal course and it is against this background that norms can sensibly lay a claim to controlling our behaviour, linking normative solutions (consequences) to kinds of possible cases (antecedents). Against this background of normalcy, true exceptions remain possible but they are, necessarily, exceptional.

It is often said that norm-givers cannot foresee the future. This is true, of course. But what puzzles me most is that they in fact *can*, albeit to a limited extent, foresee the future. Things turn out as usual, most of the time.